EXHIBIT 6

[redacted]

UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

TQ DELTA, LLC, *Plaintiff,* v. **COMMSCOPE HOLDING COMPANY, INC., COMMSCOPE INC., ARRIS US HOLDINGS, INC., ARRIS SOLUTIONS, INC., ARRIS TECHNOLOGY, INC., and ARRIS ENTERPRISES, LLC** *Defendants.*	CIV. A. NO. 2:21-CV-310-JRG (Lead Case)
TQ DELTA, LLC, *Plaintiff,* v. **NOKIA CORP., NOKIA SOLUTIONS AND NETWORKS OY, and NOKIA OF AMERICA CORP.,** *Defendants.*	CIV. A. NO. 2:21-CV-309-JRG (Member Case)

OPENING EXPERT REPORT OF BRUCE MCNAIR ON THE INVALIDITY OF THE ASSERTED CLAIMS OF THE FAMILY 6 PATENTS (U.S. PATENT NOS. 8,462,835; 8,594,162)

37. I was asked to provide an opinion regarding the skill level of a person of ordinary skill in the art of the ’835 and ’162 patents. I considered several factors, including the types of problems encountered in the art, the solutions to those problems, the pace of innovation in the field, the sophistication of the technology, my experience as a person who worked in the art on the ’835 and ’162 patents’ priority date, and the education level of active workers in the field.

38. In my opinion, at the time of the alleged invention, a person having ordinary skill in the art would have had a bachelor’s degree in electrical or computer engineering, or the equivalent, and at least 5-6 years of experience in telecommunications or a related field; a Master’s degree in electrical or computer engineering, or the equivalent, and at least 2-3 years of experience in telecommunications or a related field; or a Ph.D. in electrical or computer engineering, or the equivalent, with at least 1-2 years of experience in telecommunications or a related field.

39. I am qualified as a person of at least ordinary skill in the art, and my qualifications enable me to provide opinions regarding the ’835 and ’162 patents from the perspective of one of ordinary skill in the art.

VII. BACKGROUND OF THE TECHNOLOGY

40. Although the claims of the ’835 and ’162 patents are not limited to DSL, the applicant’s presentation of the material in the patent is in the context of asymmetric digital subscriber lines (ADSL). *See, e.g.*, ’835 patent, 1:60-2-25, 4:64-67. Therefore, I provide below a discussion of relevant DSL principles and technology.

41. The *downstream* direction is defined in DSL as the direction from the service provider to the subscriber (e.g., from the central office (CO) to the customer’s home or office). The *upstream* direction is the direction from the subscriber toward the service provider. In

ADSL, the transceiver at the subscriber's location, called the "ATU-R," transmits data in the upstream direction and receives data in the downstream direction. Conversely, the ADSL transceiver in the CO (or other location in the service provider's network), called the "ATU-C," transmits data in the downstream direction and receives data in the upstream direction. In VDSL, the transceiver at the subscriber's location, called the "VTU-R," transmits data in the upstream direction and receives data in the downstream direction, and the VDSL transceiver at the service provider's end of the subscriber line, called the "VTU-O," transmits data in the downstream direction and receives data in the upstream direction.

A. Discrete Multitone (DMT) Modulation

42. In digital communication, *modulation* is the process of generating a transmissible signal from data bits to be transmitted. Two classes of modulation are single-carrier modulation and multi-carrier modulation.

43. A single-carrier transmitter modulates one carrier signal based on the bits to be transmitted. The transmitter operates sequentially on groups of one or more bits in a bit stream to be transmitted. During each symbol period, the transmitter collects the same number of bits from the bit stream and sends a single waveform that represents the collected bits. Because the transmitter modulates a single carrier signal to transmit the data, this form of modulation is known as *single-carrier modulation*.

44. In contrast, in *multicarrier modulation*, a transmitter simultaneously modulates two or more carrier signals, known as *subcarriers* or *tones*, to transmit data to a receiver. Each subcarrier occupies and is centered within a different *subchannel*, which is a portion of the available bandwidth of the channel.

45. Typically, a multicarrier transmitter modulates a large number of subcarriers, which results in several benefits. For example, the process performed by the receiver to recover

the transmitted data can be much simpler in a well-designed multicarrier system than in a single-carrier system. Furthermore, a multicarrier transmitter can simply turn off subcarriers that would otherwise occupy subchannels that are too noisy or too attenuated to be useful. Properly designed, a multicarrier transmitter can tailor the allocation of bits to subcarriers, referred to as the *bit allocation*, *bit loading*, or *bit distribution*, in a way that is nearly optimal for the channel conditions.

46. One particular kind of multicarrier modulation, and the kind used in all ADSL and VDSL standards, is known as *discrete multitone* (DMT) modulation. DMT transmitters modulate "virtual" subcarriers by performing a mathematical operation, called an *inverse discrete Fourier transform* (IDFT), on a set of data to be transmitted. The subcarriers are inherent in the IDFT operation itself. In essence, the IDFT both creates and modulates all of the subcarriers in a single step.

47. The *discrete Fourier transform* (DFT) and IDFT are the discrete forms of, respectively, the Fourier and inverse Fourier transforms that are well known in engineering and physics. The DFT converts time-domain data into frequency-domain data, and the IDFT converts frequency-domain data into time-domain data. The DFT is often implemented using a fast Fourier transform (FFT), and the IDFT is often implemented using an inverse fast Fourier transform (IFFT).

48. As a result of the IDFT, each DMT subcarrier is modulated by a respective quadrature amplitude modulation (QAM) constellation point that represents a portion of the bit stream to be transmitted. The number of bits modulated onto each subcarrier, and therefore the bit allocation, is dependent on the signal-to-noise ratio (SNR) of the subcarrier. The "signal" part of a subcarrier's SNR is the power level of the subcarrier when it arrives at the receiver.

The power level of the subcarrier when it arrives at the receiver depends on (a) the power level at which the transmitter on the other end of the subscriber line transmitted the subcarrier and (b) the attenuation the subscriber line causes at the subcarrier's frequency. The "noise" part of a subcarrier's SNR is the power of whatever (undesirable) noise the receiver detects along with the desired subcarrier. Higher noise power levels result in lower subcarrier SNRs. Subcarriers in subchannels having higher SNRs typically carry more bits than subcarriers in subchannels having lower SNRs. Subcarriers in subchannels with SNRs that are too low to support a minimum number of bits while ensuring that the receiver can meet a target error probability do not carry any data.

49. For each DMT symbol period, the IDFT in a DMT transmitter converts discrete frequency-domain information—namely, the set of constellation points selected for the subcarriers based on the number and values of bits to be transmitted over each subchannel during the symbol period—into time-domain information that represents the amplitude and phase of the corresponding signal in time during the symbol period. These samples are referred to as *time-domain samples*. The transmitter then converts the time-domain samples to an analog signal, which it transmits to the receiver.

50. The receiver samples the analog received signal to generate a continuous stream of time-domain samples. For each symbol period, the multicarrier receiver performs a DFT on the time-domain samples in the symbol period. The DFT converts the received time-domain samples back to the frequency domain. The result is a set of rotated, attenuated, and noisy versions of the transmitted constellation points (i.e., the received points almost certainly will not coincide with the transmitted points). Assuming the number of subchannels is sufficiently

large, and the noise is not too large, the receiver recovers the original constellation points by performing a simple mathematical operation known as frequency-domain equalization.

51. Noise is an impediment to communication in both directions over subscriber lines, which emanate from the CO (or other location in the service provider's network) in cables that carry many twisted pairs in close physical proximity to each other. Although the twisting of the individual pairs within the cable helps to reduce interference (known as crosstalk) between lines, interference is unavoidable at the frequencies used in DSL, and it degrades the signal received by a receiver. In addition, the cables can allow noise from outside the cable (e.g., amateur radio transmissions) to couple into the subscriber lines. Components of the receiver also add noise, although typically at much lower levels than the noise introduced by the channel from the CO to the subscriber. If severe enough, the noise can cause the receiver to be unable to recover information sent by the transmitter. In other words, the receiver can make detection errors as it tries to reconstruct the transmitted information.

52. One particular type of noise is called *impulse noise*. Impulse noise consists of intermittent and (typically) unpredictable bursts of noise that can temporarily overwhelm the data-carrying signal for a period of time. In DSL, impulse noise is caused by "temporary electromagnetic events in the vicinity of phone lines," due to, for example, the opening of a refrigerator door, control voltages to elevators, and ringing of phones on subscriber lines sharing the same cable binder as the DSL. T. Starr, J.M. Cioffi, and P.J. Silverman, "Understanding Digital Subscriber Line Technology," Prentice-Hall (1999), p. 94. Compared to received DSL signal levels, impulse noise levels can be "enormous." *Id.* A single impulse typically lasts tens of microseconds to hundreds of microseconds, but it can last longer (e.g., 3 ms). *Id.*

53. In order to improve the possibility that the receiver will be able to recover the information from the received signal, particularly in the presence of impulse noise, DSL transceivers can use either or both of two well-known techniques that have been specified for ADSL since the release of the first ADSL standard, T1.413 Issue 1, in 1995: forward error correction (FEC) coding and interleaving, both of which I describe further below.

B. Forward Error Correction Coding

54. Error correction techniques enable a receiver to detect the presence of errors in a decoded signal and to correct at least some of those errors. To apply an error correction technique, the transmitter adds redundant information to a message to be transmitted. The receiver then uses the redundant information to detect a limited number of errors that may occur anywhere in the message, and often to correct those errors without retransmission. When the communication system uses an error correction technique, the receiver can correct at least some errors in the received data, which reduces and could eliminate entirely data lost to errors as well as any need for the receiver to request retransmission.

55. A simple example of a basic error correction code is a repetition code. The transmitter sends each message multiple times (usually an odd number of times, such as three), and the receiver chooses the message received most often as the correct message. The reconstructed message is, therefore, what the receiver determines is most likely to be the transmitted message.

56. As an example, suppose each message is 1 bit long, and the transmitter transmits each bit three times. The receiver then interprets each set of three received bits according to the following table:

Received bits	Interpreted as
000	0
001	0

010	0
100	0
011	1
101	1
110	1
111	1

57. Thus, as long as no more than one bit of the three is in error, this repetition code corrects the error and properly decodes the transmitted bit. This simple illustration is not a practical solution. It requires 2 bits of overhead data for each data bit sent, and is easily defeated by a burst of noise, such as an impulse, that spans more than a single bit. It nevertheless serves to illustrate the basic concept of error correcting codes.

58. Practical error correction techniques are often called forward error correction (FEC). There are two categories of FEC codes: convolutional codes and block codes.

59. To apply a convolutional code, the transmitter applies a convolutional encoding algorithm that operates at the bit level. The convolutional code converts the data to be transmitted into coded data that includes redundancy for error detection and correction. A well-known example of convolutional coding is the trellis coding that can be used in ADSL. Trellis coding is not relevant to the ’835 and ’162 patents and is not discussed further in this report.

60. To apply a block code, the transmitter partitions the data to be transmitted into blocks and adds redundancy to each block to form a *codeword*. The transmitter appends r redundancy elements to each block of k elements, without altering the original k elements, to form codewords having k + r elements. The receiver then uses the information in the r redundancy elements to automatically detect and correct errors in the data portion of the codeword. There are many types of FEC block codes. For example, a parity code is a type of

FEC block code in which the encoder appends a bit to the end of a group of bits, where the appended bit indicates whether the number of "1s" in the group of bits is even or odd.

61. Another well-known example of a FEC block code is the Reed-Solomon code that has been used in ADSL since the mid-1990s. Reed-Solomon codes encode blocks of data to be transmitted in such a way that errors in certain elements of the block can be corrected as long as enough other elements of the block are error-free. As long as enough elements of a block are error-free, the redundancy symbols allow the receiver to correct all of the errors in the rest of the block.

62. The N elements in each Reed-Solomon codeword include K data elements (i.e., the data to be transmitted in the codeword) and R redundancy elements that result from the Reed-Solomon code calculation. Thus,

$$N = K + R.$$

63. The mathematics of Reed-Solomon codes are complicated, but the essential property is that the maximum number of data elements with any number of errors that can be corrected is equal to half of the number of redundancy elements. In other words, a Reed-Solomon code can correct as many as R/2 errored data elements, regardless of how many individual errors are within those data elements.[5]

64. As a simple example, if the elements in a Reed-Solomon codeword are bytes, and R = 16, up to 8 erroneous bytes can be corrected, regardless of how many bits of each errored byte are in error. Whether a byte contains one bit error or eight, the Reed-Solomon code will correct all bit errors in up to 8 bytes. Thus, Reed-Solomon codes are particularly helpful

[5] This discussion assumes that erasures are not used.

for use in environments that suffer from impulse noise because any erroneous element counts as a single error regardless of how many bits of that element are in error.

65. Figure 1 below is a block diagram showing a Reed-Solomon encoder in the transmitter for a byte-oriented system:

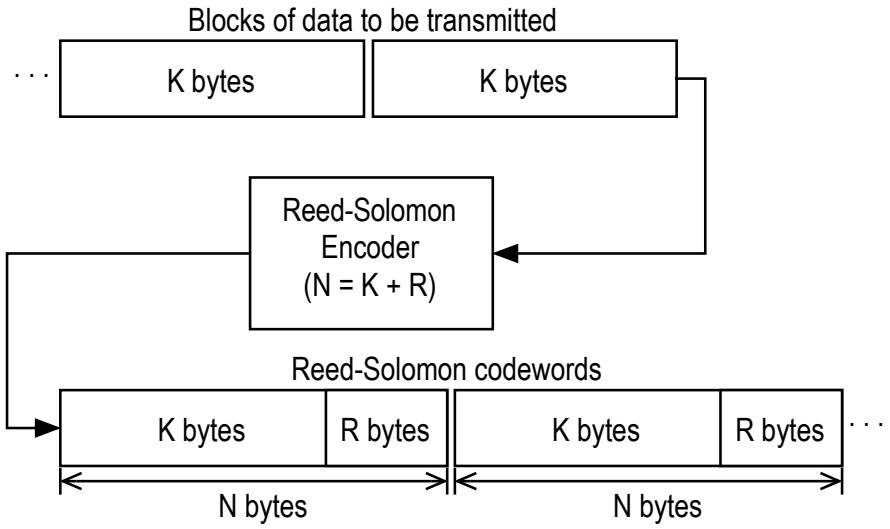

Figure 1: Reed-Solomon encoder for byte-oriented system

66. Figure 2 illustrates a Reed-Solomon decoder in the receiver of the same system:

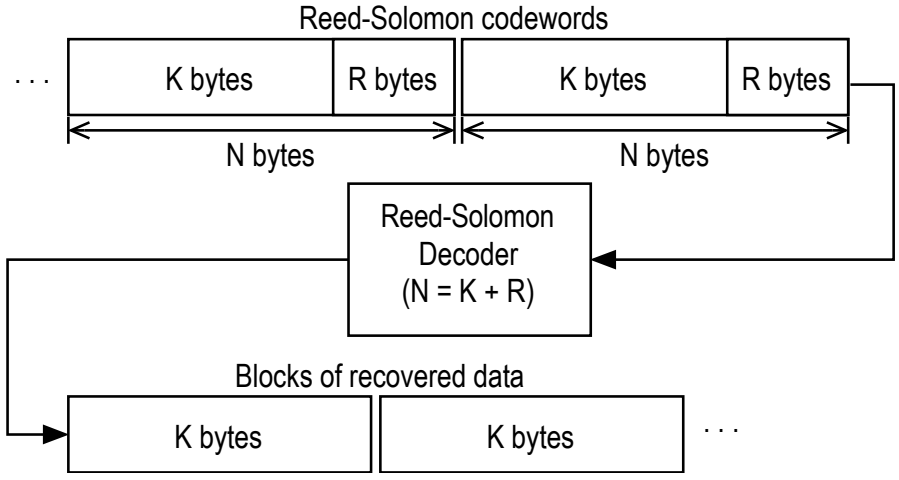

Figure 2: Reed-Solomon decoder for byte-oriented system

67. As explained above, the power of a Reed-Solomon code is directly proportional to the number of redundancy bytes, R. Codes with larger values of R can correct more errored bytes.

68. In order to generate the redundancy bytes and the Reed-Solomon codeword, the transmitter must first collect all of the data bytes for that codeword. Similarly, in order to detect and correct errors in the received data, the receiver must wait until it has received the entire Reed-Solomon codeword. Thus, the use of FEC such as Reed-Solomon coding delays the processing of data in both the transmitter and the receiver.

69. The redundancy elements of a FEC code must be transmitted in addition to the user data and thus constitute overhead that reduces the amount of user data that can be transmitted per unit of time. Therefore, designers of communication systems using FEC select the coding parameters (e.g., the amount of redundancy) to provide an acceptable trade-off between error detection/correction capability and overhead.

C. Interleaving

70. Interleaving is a technique used in many digital communication systems to improve the effectiveness of FEC block codes, such as Reed-Solomon codes, when the transmitted signal is affected by impulse noise. Like FEC, interleaving has been specified for use in ADSL since the first standard, T1.413 Issue 1, was finalized in 1995. To apply interleaving, the transmitter shuffles a certain number of consecutive data elements (e.g., bytes) in a known and systematic way before transmission over the channel. As a result of the interleaving process, data elements that are adjacent in the data sequence are transmitted non-contiguously, spread out over a time interval. After decoding the data elements, the receiver reassembles them in order prior to delivery to their ultimate destination.

71. Figure 3 below illustrates how the interleaving process works using a simple interleaving procedure known as block interleaving. The transmitter separates the data elements into a number of blocks of five elements each (i.e., A1, A2, A3, A4, A5; B1, B2, B3, B4, B5; etc.). Elements from four blocks are then interleaved and transmitted in turn to create a sequence AxBxCxDx.

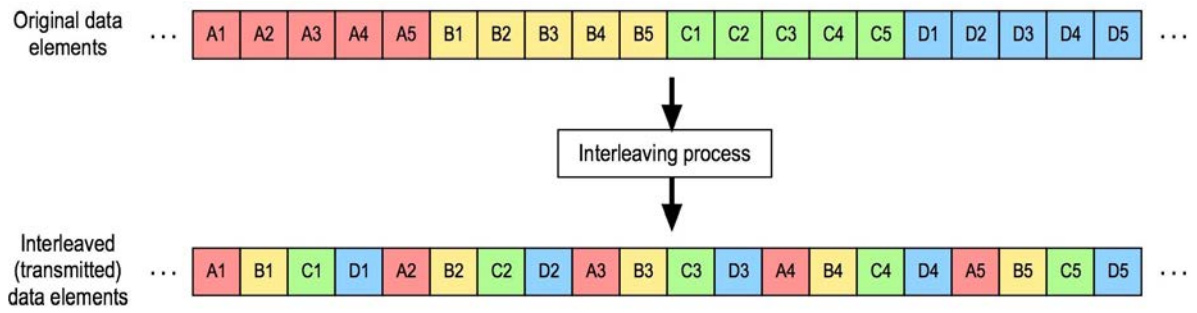

Figure 3: Interleaving example

72. Note that the interleaving process significantly increases the time interval over which the data elements from each of the individual data blocks are transmitted. Following the interleaving process in this simple example, the first and last data elements of each block are separated by 15 other data elements, and consecutive data elements of each block are separated by three other data elements. In real systems, the separation after interleaving will normally be hundreds or even thousands of bytes.

73. The receiver knows how the transmitter interleaved the data elements and can reverse the interleaving process by collecting all of the interleaved data elements and applying a complementary deinterleaving process. The deinterleaving process results in the data elements being reordered into their original sequence, as shown in Figure 4.

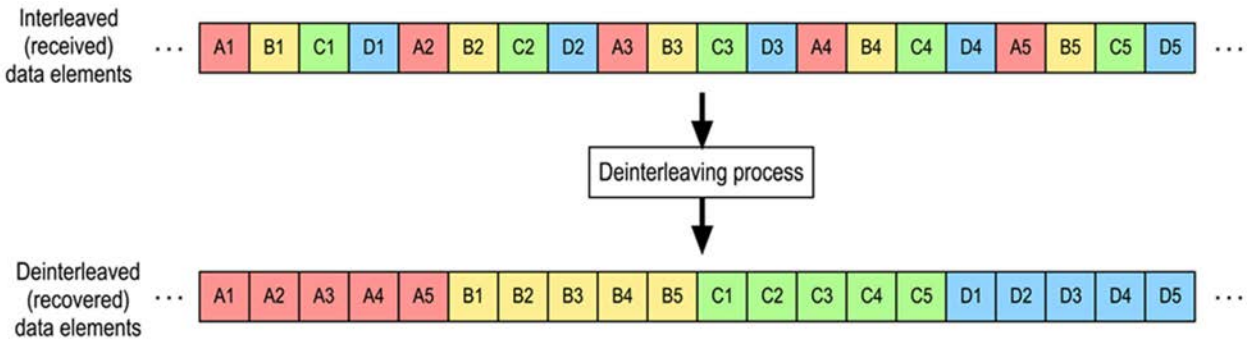

Figure 4: Deinterleaving example

74. Digital communication systems often use interleaving as a strategy to improve the performance of FEC block codes. Errors in the received data tend to occur in bursts (e.g., as a result of impulse noise). As explained above, a Reed-Solomon code can correct up to R/2 errored data elements in a block. If a single block of the received data has more than R/2 data elements with errors, the Reed-Solomon code cannot correct all of the errors. If the transmitter interleaves the data elements prior to transmission, however, the effects of bursts of errors are spread out in time and over multiple Reed-Solomon codewords when the data elements are reordered by the receiver.

75. An example illustrates the potential benefits of a communication system using interleaving with FEC. Suppose the transmitter separates the data elements into blocks of five elements each and applies Reed-Solomon FEC with R = 4 to the blocks, but the transmitter does not use interleaving. Figure 5 below shows the effect of an occurrence of impulse noise on the transmitted data elements. As a result of the impulse noise, the received data elements B5, C1, C2, C3, and C4 have errors. The Reed-Solomon code can correct up to two data elements within each block, which corrects the errors in the data element B5. The Reed-Solomon code cannot correct the errors in all of the data elements C1, C2, C3, or C4, however,

because they are all within the same block, and the number of errored data elements in the block (4) exceeds the error-correcting capability of the Reed-Solomon code (R/2 = 2).

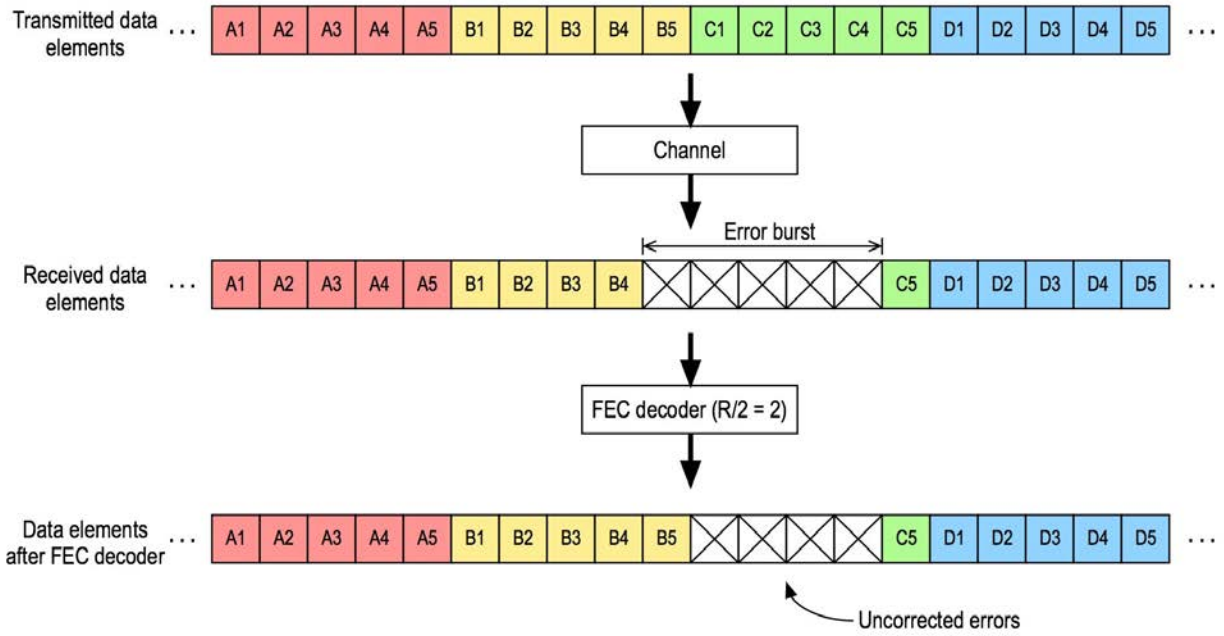

Figure 5: Effect of error burst caused by impulse noise without interleaving

76. Figure 6 below shows how the use of interleaving can spread out the effects of impulse noise over time and over multiple Reed-Solomon codewords, thereby increasing the probability of the Reed-Solomon code being able to correct all of the errored data elements in each of the individual codewords. Because the transmitter interleaved the elements of the codewords prior to transmission, the impulse noise that caused four uncorrectable errors in the scenario shown in Figure 5 has no effect on the FEC decoded data in the example of Figure 6. Following the deinterleaving process in the receiver, none of the blocks of data elements has more than two data elements with errors. And because the Reed-Solomon code can correct up to two data elements per block, the Reed-Solomon code is able to correct all of the errors caused by the impulse noise.

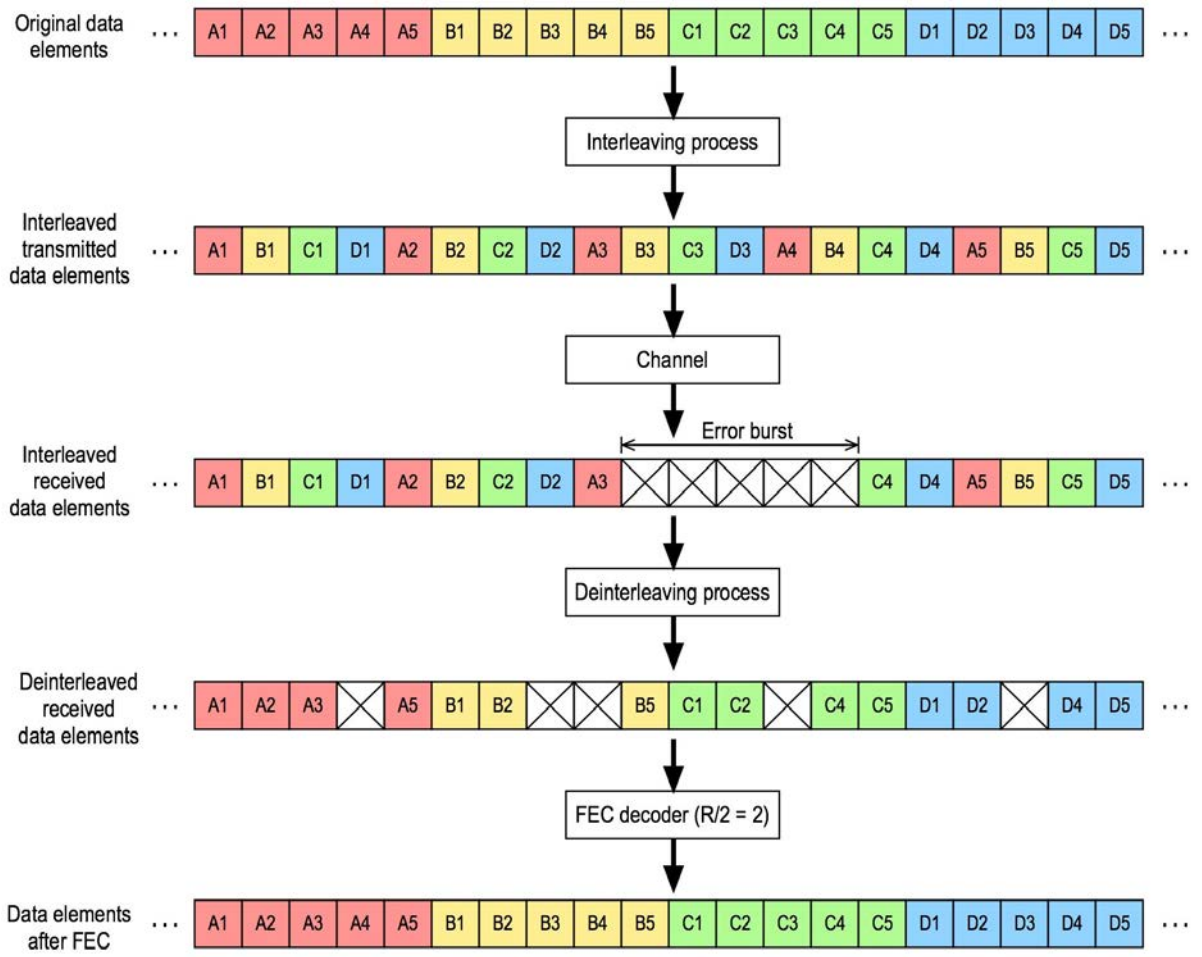

Figure 6: Effect of interleaving with burst errors caused by impulse noise

77. As illustrated by the example above, interleaving can improve the effectiveness of FEC block codes by spreading the effects of bursts of errors over time. The more spread out the formerly-consecutive data elements are after interleaving, the more the interleaving can improve the performance of the FEC block code to mitigate the effects of impulse noise. The *interleaver depth* is the minimum separation after interleaving between any two data elements that were adjacent before interleaving. In the example in Figure 6 above, the interleaver depth is 4.

78. The process of interleaving in the transmitter and deinterleaving in the receiver introduces delay (or latency) because bytes are stored for a period of time after they are written

to the interleaver or deinterleaver memory but before they are read out. In units of time, the end-to-end interleaver delay depends on the rate at which the transmitter transmits and the receiver receives bytes, i.e., the bit rate over the communication channel, because that rate determines how quickly the transceivers process the bytes. For the type of interleaver specified in ADSL and VDSL, the end-to-end interleaver delay, in milliseconds, is given by the equation $(I - 1) \times (D - 1) \times 8/Rate$, where I is the interleaver block length, D is the interleaver depth in bytes, and *Rate* is the bit rate over the line in kbit/s. Thus, the latency (delay) in units of time is directly proportional to the interleaver depth and inversely proportional to the line rate.

79. Because the latency due to interleaving increases with the interleaver depth, there is a fundamental trade-off between latency and improved FEC block code performance. A larger interleaver depth offers better potential to improve the performance of FEC block codes but also increases latency.

80. Different types of data have different levels of sensitivity to delays that occur during transmission. For example, real-time voice data, such as from a telephone call, cannot tolerate large delays because delays in conversations are very annoying for the participants. In contrast, a pre-stored (i.e., non-real-time) video transmitted to a receiver (e.g., via YouTube, Netflix, HBO Go, etc.) can tolerate a substantial amount of delay because the local video playback device typically buffers at least some of the video before playing it. Therefore, in DSL systems, the maximum allowed latency, in ms, is specified as a constraint. The transceivers then configure their FEC and interleaving parameters so that the latency, in ms, at the bit rate transmitted over the line, is less than or equal to the maximum allowed latency.

D. Overview of DSL Standards

81. Many of the aspects of DSL have been standardized. These aspects include the use of Reed-Solomon coding and interleaving to mitigate errors, and procedures to allow two connected transceivers to change their configurations and settings to accommodate changes in the transmission environment or the type of data to be transmitted.

82. Every DSL standard defines procedures to enable the transceivers on either end of a subscriber line to establish a connection and procedures that the transceivers conduct to maintain an established connection. The state in which the DSL transceivers establish a connection is referred to as *initialization*, and the state that immediately follows initialization, during which the transceivers can transfer user data, is called *Showtime*.

83. Multiple standardization organizations have been involved in sometimes-overlapping work. This section introduces two standardization organizations whose work predated the '835 and '162 patents, explains their relationship, and introduces the standards that are particularly relevant to the subject matter of this case and pre-date the priority date of the '835 and '162 patents.

84. The Alliance for Telecommunication Industry Solutions (ATIS) had a large role in the early standardization of ADSL. ATIS is a standardization organization based in the United States and accredited by the American National Standards Institute (ANSI). Generally, ATIS seeks to develop technical standards that ensure interoperable end-to-end telecommunication solutions that can be timely implemented.

85. ATIS was working to standardize DSL by the early 1990s. The T1E1.4 working group of ATIS carried out the earliest work on ADSL standardization, and, in 1995, ATIS published the ANSI ADSL standard T1.413, which I will refer to herein as "T1.413 Issue 1."

The T1E1.4 working group continued its work on ADSL after 1995, and, in 1998, ATIS published a revision of T1.413 Issue 1, which I will refer to as "T1.413 Issue 2."

86. The other standardization organization relevant to this case is the International Telecommunication Union (ITU). The ITU is an agency of the United Nations that specializes in information and communication technologies. The Telecommunications Sector of the ITU (known as the ITU-T) defines transceiver specifications for use internationally. The ITU-T generates standards called "Recommendations" for all fields of telecommunications. In relation to DSL, the Recommendations are primarily aimed at defining mandatory and optional functions of DSL transmitters.

87. Standardization work in the ITU-T is carried out by Study Groups, and DSL standardization was and is carried out by Study Group 15. The work of Study Group 15 is partitioned into work areas known as "Questions." DSL standardization work takes place within Question 4 (abbreviated herein as SG15/Q4).

88. In 1997, SG15/Q4 began working on ADSL standardization. Among other projects, SG15/Q4 established the projects known as G.dmt and G.lite. At that time, T1.413 Issue 1 was in force, and T1E1.4 was developing what eventually became T1.413 Issue 2. G.dmt was expected essentially to adopt T1.413 Issue 2 and add annexes addressing country-specific issues. The work in G.dmt was eventually released in 1999 as ITU-T Recommendation G.992.1, much of which is identical to T1.413 Issue 2. G.lite was expected to be a lower-speed version of ADSL. The work in G.lite was eventually released, also in 1999, as ITU-T Recommendation G.992.2.

89. Once T1E1.4 completed T1.413 Issue 2, SG15/Q4 became the lead standards working group responsible for defining DSL transceiver standards. After releasing T1.413 Issue

2, T1E1.4 focused less on transceiver standards and more on providing inputs to SG15/Q4 regarding North American requirements and preferences for emerging DSL standards.

90. By 2004, the in-force ITU-T Recommendations relevant to ADSL, and relevant to this case, included G.992.1, G.992.2, G.992.3, and G.992.4. Each of these Recommendations specifies techniques for transmitting a range of bit rates over the existing copper local network from high bit rates over relatively short distances to lower bit rates over longer distances. Like T1.413 Issue 1 and T1.413 Issue 2, all of the ITU-T's ADSL Recommendations specify the use of DMT modulation.

91. The Recommendation G.992.1, entitled "Asymmetric Digital Subscriber Line (ADSL) Transceivers" and sometimes referred to by those skilled in the art as "ADSL1" or "G.dmt," specifies the physical layer characteristics of an ADSL transceiver at the subscriber line interface. G.992.1 specifies transceiver requirements for both the transceiver at the telephone company's central office (the ATU-C) and the transceiver at the subscriber's premises (the ATU-R) to enable connections that support at least 6.144 Mbit/s in the downstream direction (i.e., toward the subscriber) and at least 640 kbit/s in the upstream direction (i.e., away from the subscriber) without interfering with plain old telephone signals (POTS) on the same subscriber line.

92. Like G.992.1, Recommendation G.992.2, entitled "Splitterless Asymmetric Digital Subscriber Line (ADSL) Transceivers" and sometimes referred to as "G.lite," specifies the physical layer characteristics of an ADSL transceiver at the subscriber line interface. But unlike G.992.1, which was designed to maximize performance, G.992.2 was designed to provide lower bit rates (i.e., a maximum of 1.536 Mbit/s downstream and 512 kbit/s upstream)

by simplifying various aspects of the transceiver, such as not requiring a splitter to separate the DSL signals from POTS signals on the subscriber line.

93. Recommendation G.992.3, entitled "Asymmetric Digital Subscriber Line Transceivers 2 (ADSL2)" and sometimes referred to as "ADSL2" or, during its development, as "G.dmt.bis," builds on many of the aspects of G.992.1 to enable connections that support at least 8 Mbit/s downstream and at least 800 kbit/s upstream. G.992.3 also adds a number of features and capabilities, some of which are discussed later in this report.

94. Recommendation G.992.4, entitled "Splitterless asymmetric digital subscriber line transceivers 2 (splitterless ADSL2)," referred to as "G.lite.bis" during its development, builds on G.992.2, adding a number of features and capabilities.

1. T1.413 Issue 1 (1995)

95. The world's first ADSL standard, known as "T1.413 Issue 1," was released by ATIS in 1995. This section provides an overview of some of the features of T1.413 Issue 1 that are relevant to the '835 and '162 patents.

a. Initialization

96. T1.413 Issue 1 specifies an initialization procedure during which, among other things, the transmitter and receiver agree on the number of bits each subcarrier will transport, i.e., the bit allocation. *See, e.g.*, T1.413 Issue 1, § 12; *see also*, *e.g.*, T1E1.4/93-184, p. 1 ("During initialization, the communication channel is tested, and an optimal bit allocation table is determined for the desired data rate.").

b. Frames and Superframes

97. T1.413 Issue 1 organizes the bits to be transmitted into *frames*. The transmitter transmits all of the bits in a frame in the same DMT symbol, with bits allocated to subcarriers in accordance with the bit allocation determined during the initialization procedure. To

construct the DMT symbol, the bits in each frame are separated into groups of bits, each group being allocated to a different subcarrier. The number of bits in each group may be different.

98. T1.413 Issue 1 defines a collection of 69 consecutive frames as a *superframe*. Figure 7 below shows the superframe structure. Each superframe is composed of 68 data frames, numbered 0 through 67, followed by a symbol known as the *synchronization symbol* (or "synch symbol," sometimes spelled as "sync symbol"). *See, e.g.*, T1.413 Issue 1, § 6.2.1.1.

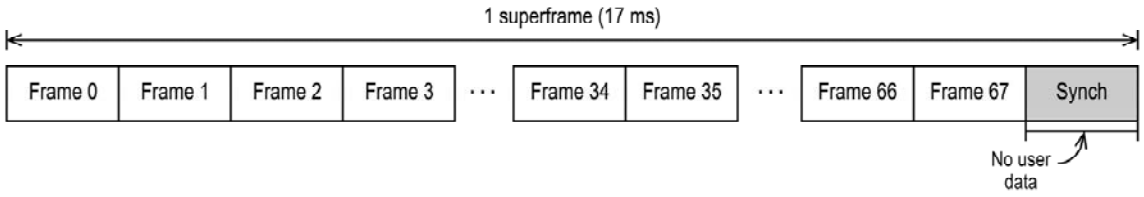

Figure 7: T1.413 Issue 1 superframe

Each data frame in a superframe carries data. In contrast, the synchronization symbol does not carry any data. Instead, it is a special symbol, known to the receiver, to help the receiver identify the superframe boundaries (and thus the frame boundaries). *See, e.g.*, T1.413 Issue 1, § 6.9.1.2 ("Therefore a gross timing error that is an integer multiple of 8 samples could still persist after a micro-interruption (e.g., a temporary short-circuit, open circuit or severe line hit); correction of such timing errors is made possible by the use of the synchronization symbol defined in 6.9.3."); *id.* at § 6.9.3 ("The synchronization symbol permits recovery of the frame boundary after micro-interruptions that might otherwise force retraining."); *id.* at § 7.9.3 (same). The synchronization symbol is the same during every superframe. *Id.* at § 6.9.3 ("a synchronization symbol (with a nominal length of 544 samples) inserted after every 68 data symbols."); *id.* at § 7.9.3 (a synchronization symbol (with nominal length of 68 samples) is inserted after every 68 data symbols).

99. The bits used to generate the synchronization symbol are from a pseudo-random sequence. *See, e.g.*, T1.413 Issue 1, § 6.9.3 ("The data pattern used in the synchronization symbol shall be the pseudo-random sequence PRD, (dn, for n = 1 to 512) defined by . . . ."); *id.* at § 7.9.3 ("The data pattern used in the-synchronization symbol is the pseudo-random sequence PRU (dn, for n = 1 to 64), defined by . . . ."). The pseudo-random sequence used to generate the synchronization symbol is the same pseudo-random sequence that is used during the initialization procedure to generate the REVERB and MEDLEY signals. *See, e.g.*, *id.* at § 12.4.4 ("The data pattern used in C-REVERB1 shall be the pseudo-random downstream sequence (PRD), dn for n = 1 to 512, defined in 6.9.3 and repeated here for convenience. . . ."); *id.* at § 12.5.2 ("The data pattern used in R-REVERB1 shall be the pseudo-random upstream sequence PRU defined in 7.9.3 and repeated here for convenience. . . ."); *id.* at § 12.6.6 ("C-MEDLEY is a wideband pseudo-random signal used for estimation at the ATU-R of the downstream SNR. The data to be transmitted shall be derived from the pseudo-random sequence, PRD, and modulated as defined in 6.9.3 and 12.4.4."); *id.* at § 12.7.8 ("R-MEDLEY is a wideband pseudo-random signal used for estimation of the upstream SNR at the ATU-C. The data to be transmitted are derived from the pseudo-random sequence PRU defined in 12.5.2, continuing from one symbol to the next."). The reuse of the same pseudo-random sequence for multiple purposes reduces the complexity of the transceivers.

c. Latency Paths, FEC, and Interleaving

100. As explained above, it is well known that different types of data have different levels of sensitivity to delays that occur during transmission. T1.413 Issue 1 was designed to allow a single ADSL connection to transfer multiple types of data having different latency requirements (or delay tolerances). *See, e.g.*, T1E1.4/93-117, p. 2 ("We maintain the two-buffer

parameterization explained in [3] that separates the data into delay-tolerant and delay-intolerant streams.").

101. To allow the transmission of both delay-sensitive and delay-tolerant data over the same subscriber line, T1.413 Issue 1 supports *dual-latency* configurations using two independent *data paths*. Delay-sensitive data is routed to the so-called *fast path*, and delay-tolerant data is routed to the *interleaved path*. To mitigate errors in data transmission and reception, T1.413 Issue 1 specifies the use of FEC in the fast path, and both FEC and interleaving in the interleaved path. *See, e.g.*, T1.413 Issue 1, § 6.4.1 (Reed-Solomon coding), § 6.4.2 (interleaving). Thus, the key difference between the fast and interleaved paths is that data in the interleaved path is both FEC coded and interleaved, as described previously, to improve the effectiveness of the FEC code in the presence of impulse noise. In contrast, the fast path includes FEC coding, but does not include interleaving. Therefore, the fast path has lower latency than the interleaved path. *See, e.g.*, T1.413 Issue 1, §§ 6.4.1, 6.4.2.

d. <u>On-Line Reconfiguration: Bit Swapping</u>

102. As explained above, the number of bits that can be modulated onto each DMT subcarrier is dependent on and proportional to the SNR of that subcarrier. Because the number of bits each subcarrier can support is dependent on its SNR, and because SNR is a receiver quantity, T1.413 Issue 1 systems undergo an initialization procedure that, among other things, allows the receiver to estimate the SNR of each subcarrier on which it can potentially receive data. The initialization procedure is followed by what is known in the DSL art as "Showtime." Showtime is the phase in which steady-state communication takes place.

103. The receiver estimates the subcarrier SNRs during initialization, and it determines the bit loading. The receiving transceiver then communicates the bit loading to the transmitting transceiver. The transmitter uses the bit loading to modulate bits onto the

subcarriers, and the receiver uses the bit loading to recover the bits from the received subcarriers. The transmitter and receiver must use the same bit loading at all times. If they do not, the receiver will likely make detection errors, which could cause the connection to fail.

104. When Showtime begins, the transmitter and receiver use the bit loading determined and communicated during initialization. As Showtime progresses, the subcarrier SNRs can change, typically slowly, and typically due to changes in the noise detected by the receiver, but potentially also due to changes in the channel. If the SNR of a subcarrier degrades over time to the point that the number of bits modulated onto it cannot be reliably decoded by the receiver, the receiver will make detection errors, which, if severe enough, can cause the connection to fail.

105. During the development of T1.413 Issue 1, it was recognized that it would be desirable to adaptively modify "the bit allocation table and transmission bandwidth of an ADSL signal . . . in response to changing channel conditions," such as "the introduction of additional cross-talkers and the long term modification of the channel impulse response due to temperature, weather, or other changing line conditions." T1E1.4/93-184, p. 3. Participants in T1E1.4 recognized that "[t]he ability to adapt to these changes without interrupting data flow will greatly enhance the performance of the ADSL system." *Id.*

106. Accordingly, to decrease the likelihood of connection failures during Showtime, T1.413 Issue 1 defines a procedure for "[o]n-line adaptation and reconfiguration" called *bit swapping*, which "enables an ADSL system to change the number of bits assigned to a subcarrier, or change the transmit energy of a subcarrier without interrupting data flow." T1.413 Issue 1, §§ 13, 13.2. Using bit swap procedures, the receiver can instruct the transmitter on the other side of the subscriber line to make modest changes to the bit loading and transmit power

allocation to compensate for changes the receiver detects over time in the subcarrier SNRs. For example, if the receiver detects that the SNR of a subcarrier has decreased substantially and consistently over time, the receiver can instruct the transmitter to load fewer bits onto that subcarrier and to load more bits onto another subcarrier or set of subcarriers that have sufficient SNR to accommodate additional bits. Thus, for example, the receiver might decide that a subcarrier currently carrying 6 bits should instead only carry 5 bits (by changing the constellation diagram from 64-QAM to 32-QAM), and another subcarrier carrying only 3 bits can accommodate one additional bit to carry 4 bits (by changing the constellation diagram from 8-QAM to 16-QAM). Therefore, the total number of bits carried by each DMT symbol remains the same after a bit swap, but the allocation of those bits to the subcarriers changes.

107. In order to determine when a bit swap is needed, ADSL receivers monitor the subchannel SNRs during Showtime.

108. The bit swap procedure in T1.413 Issue 1 uses the ADSL overhead control (aoc) channel, which is a logical channel carried in the interleaved path. T1.413 Issue 1, § 13.1. To provide robustness, "[a]ll aoc messages are transmitted 5 consecutive times for extra security," and "[a] transceiver unit shall only act on an aoc message if it has received three identical messages in a time period spanning 5 of that particular message." *Id.* at § 13.1.2. Thus, "[a]ll bit swap messages shall be repeated five consecutive times" over the aoc channel. *Id.* at § 13.2.1.

109. Bit swaps are coordinated using superframe counters. T1.413 Issue 1, § 13.2.2. The ATU-C and ATU-R transmitters and receivers set their superframe counters to 0 at the beginning of Showtime. *Id.* After transmitting each superframe, each transmitter increments its counter, and each receiver increments its counter after receiving each superframe. *Id.* The

transmitter and receiver superframe counters are kept synchronized during Showtime "using the synch symbol in the ADSL frame structure." *Id.*

110. To conduct a bit swap, the receiving transceiver transmits a bit swap request to the transmitting transceiver. T1.413 Issue 1, § 13.2.3. The bit swap request tells the transmitter which subcarriers need to be modified, and how to modify them. *Id.* To provide robustness, the receiving transmitter transmits the bit swap request message five consecutive times. *Id.*; *see also* T1E1.4/93-184 ("For data integrity, any bit swap message will be repeated five times. Note that this is in addition to any trellis coding or forward error correction already in the system. A unit will act on the message only if it receives 3 or more identical copies. For a channel with a BER of 1e-7, this gives the probability of receiving an erroneous bit swap message at less than 1e-14.").

111. After the transmitting transceiver receives three identical bit swap request messages within a set period of time, it transmits a bit swap acknowledge message to the receiving transmitter. T1.413 Issue 1, § 13.2.5. The bit swap acknowledge message includes the "superframe counter number, which indicates when the bit swap is to take place." *Id.* Like the bit swap request message, the bit swap acknowledge message is transmitted five times to provide robustness. *Id.*; *see also* T1E1.4/93-184 ("For data integrity, any bit swap message will be repeated five times. Note that this is in addition to any trellis coding or forward error correction already in the system. A unit will act on the message only if it receives 3 or more identical copies. For a channel with a BER of 1e-7, this gives the probability of receiving an erroneous bit swap message at less than 1e-14.").

112. After the receiving transceiver receives three identical bit swap acknowledge messages within a set period of time, it waits until the superframe counter value is equal to the

value specified in the bit swap acknowledge message. T1.413 Issue 1, § 13.2.6. Beginning with the first frame of the next superframe, the receiver implements the bit swap, thereby changing the assignment of bits and power to subcarriers. *Id.*

113. Likewise, after transmitting the bit swap acknowledge message five times, the transmitting transceiver waits until the superframe counter value is equal to the value it specified in the bit swap acknowledge message. T1.413 Issue 1, § 13.2.7. Beginning with the first frame of the next superframe, the transmitter also implements the bit swap, thereby changing the assignment of bits and power to subcarriers synchronously with the receiver. *Id.*

114. Thus, T1.413 Issue 1 allows bit swapping to be executed only at the boundaries between superframes.

2. T1.413 Issue 2 (1998) and ITU-T Recommendation G.992.1 (1999)

115. In 1998, ATIS published a revision of T1.413 Issue 1 ADSL, known as "T1.413 Issue 2." In 1999, the ITU released ITU-T Recommendation G.992.1, known as "ADSL1," much of which is identical to T1.413 Issue 2. Accordingly, I will describe aspects of T1.413 Issue 2 and G.992.1 together in this subsection. In general, I will refer to G.992.1; when necessary, I will point out where G.992.1 and T1.413 Issue 2 differ.

116. ITU-T Recommendation G.992.1 and T1.413 Issue 2 specify requirements for ADSL transceivers to enable high-speed data transmission between the network operator end of a subscriber line, where the ATU-C is located, and the customer end, where the ATU-R is located. G.992.1, i. The acronym "ATU" stands for "ADSL transceiver unit." *Id.* at § 4.

117. One objective of G.992.1 is to provide "compatibility and interoperability between transceivers complying to this Recommendation and between transceivers that include different combinations of options." *Id.* at i. G.992.1 defines two "categories" of ADSL transceivers: category I includes "transceivers with no performance-enhancing options, which

meet a basic set of performance requirements," and category II includes "transceivers with performance-enhancing options which meet an expanded set of performance requirements." *Id.* at §§ 3.7, 3.8.

118. Figure 1-1 of G.992.1, copied below, is an ADSL system reference model, which illustrates the locations of the ATU-C and ATU-R within the network, as well as various defined interfaces between components.

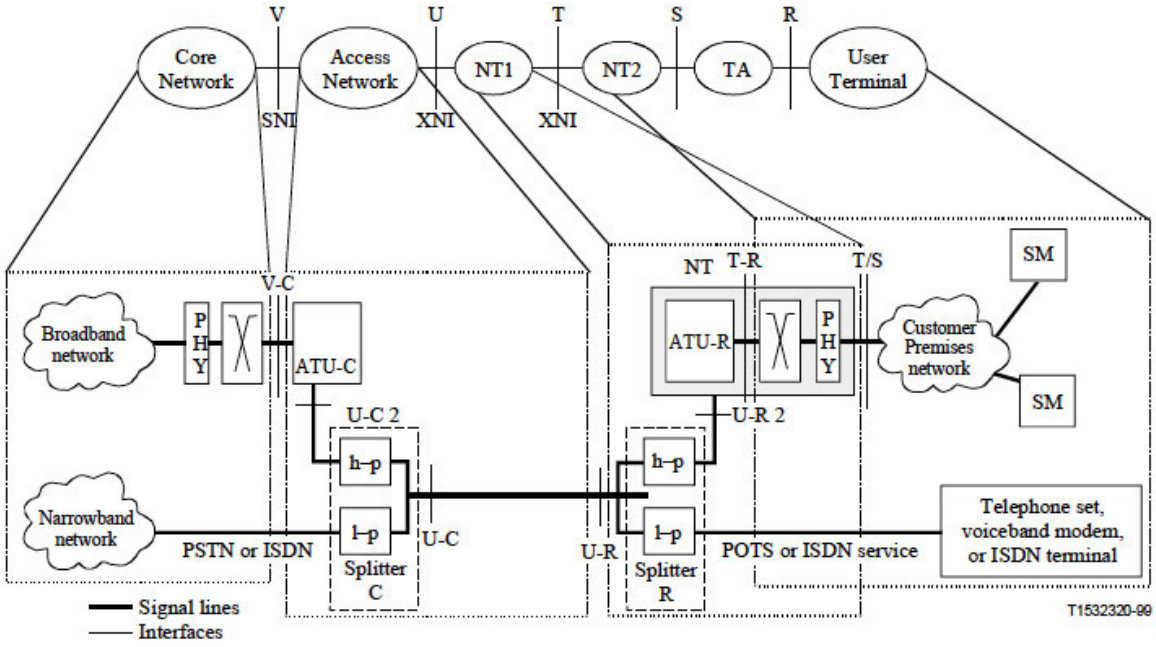

NOTE 1 – The U-C and U-R interfaces are fully defined in this Recommendation. The V-C and T-R interfaces are defined only in terms of logical functions, not physical. The T/S interface is not defined here.
NOTE 2 – The V-C interface may consist of interface(s) to one or more (STM or ATM) switching systems.
NOTE 3 – Implementation of the V-C and T-R interfaces is optional when interfacing elements are integrated into a common element.
NOTE 4 – One or other of the high-pass filters, which are part of the splitters, may be integrated into the ATU-x; if so, then the U-C 2 and U-R 2 interfaces become the same as the U-C and U-R interfaces, respectively.
NOTE 5 – A digital carrier facility (e.g. SONET extension) may be interposed at the V-C.
NOTE 6 – Due to the asymmetry of the signals on the line, the transmitted signals shall be distinctly specified at the U-R and U-C reference points.
NOTE 7 – The nature of the customer installation distribution and customer premises network (e.g. bus or star, type of media) is for further study.
NOTE 8 – More than one type of T-R interface may be defined, and more than one type of T/S interface may be provided from an ADSL NT (e.g. NT1 or NT2 types of functionalities).
NOTE 9 – A future issue of this Recommendation may deal with customer installation distribution and home network requirements.
NOTE 10 – Specifications for the splitters are given in Annex E.

Figure 1-1/G.992.1 – ADSL system reference model

119. Among other things, G.992.1 defines the modulation technique used in ADSL (namely, DMT), describes how data and overhead are organized into frames that are modulated

for transmission, defines an initialization protocol, and defines a protocol that the two transceivers can use to modify their FEC and/or interleaver parameters during steady-state transmission. *Id.* at §§ 1, 7, 8, 10; *id.* at Appendix II.

120. Figure 5-1 of G.992.1, copied below, is a block diagram of an ATU-C. *Id.* at § 5.1.1; *see also id.* at Figure 5-2. G.992.1 provides similar block diagrams for the ATU-R. *See, e.g.*, Figures 5-3, 5-4.

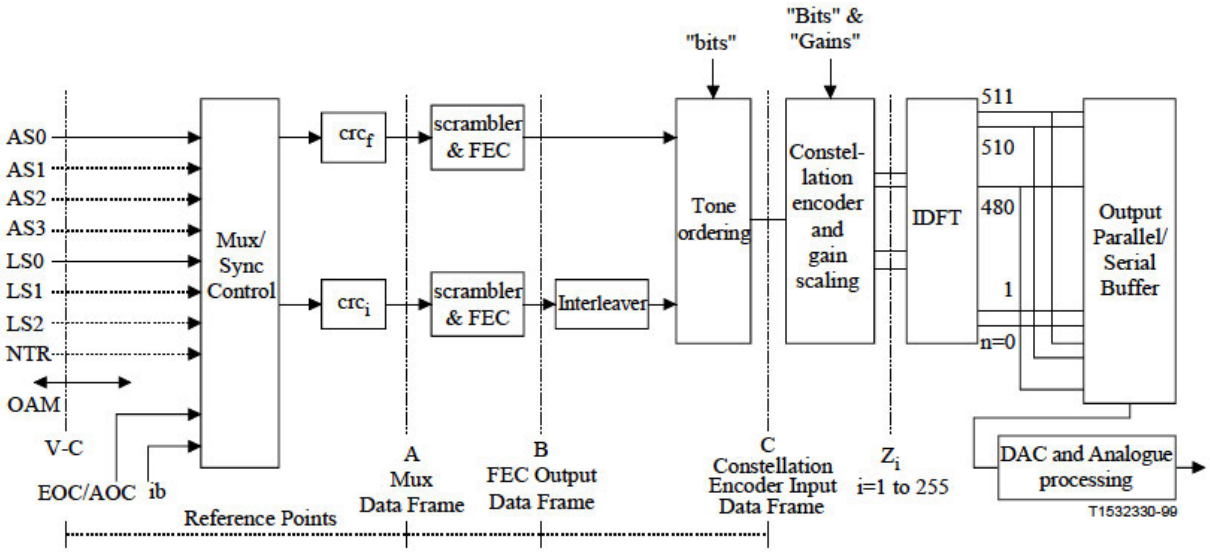

NOTE – Solid versus dashed lines are used to indicate required versus optional capabilities respectively. This figure is not intended to be complete in this respect, see clauses 6 and 7 for specific details.

Figure 5-1/G.992.1 – ATU-C transmitter reference model for STM transport

121. Because of the addition of FEC redundancy bytes (in both the fast and interleaved paths) and interleaving, both of which are discussed further below, the data frames have a different appearance at different points along the data processing path. As shown above, Figure 5-1 of G.992.1 shows three reference points A, B, and C. The reference point A follows assembly of a *mux data frame*. The reference point B follows the mux data frames being scrambled and FEC encoded. The reference point C follows interleaving for the interleaved

data, discussed below, and a process called tone ordering, which is not important for the discussion herein.

122. G.992.1 specifies the use of Reed-Solomon FEC coding for both fast and interleaved data. *See, e.g.*, *id.* at § 7.6.1 ("R . . . redundant check bytes . . . shall be appended to K . . . message bytes . . . to form a Reed-Solomon codeword of size N = K + R bytes."). G.992.1 defines a parameter, S, which specifies the number of mux data frames per FEC codeword. *See, e.g.*, *id.* at § 8.4.1.2 ("S = number of mux data frames per FEC codeword"). For the fast path, the value of S is always 1 (*see, e.g.*, *id.* at §§ 7.4.1.2.1, 7.6, Table 7-7), which means that each DMT symbol carries exactly one FEC codeword. In other words, in the fast path, the FEC codeword boundaries are identical to and coincident with the DMT symbol boundaries. For interleaved data, G.992.1 also requires the use of a convolutional interleaver with a programmable interleave depth. *Id.* at § 7.6.3 ("The Reed-Solomon codewords in the interleave buffer shall be convolutionally interleaved. The interleaving depth varies, as explained in 7.4, but shall always be a power of 2.").

123. G.992.1 sets forth minimum required downstream and upstream FEC and interleaving capabilities for both the ATU-C and ATU-R, including all of the values of R, S, and D the transceivers must be capable of supporting. *See, e.g.*, G.992.1, §§ 7.6, 8.6. Table 7-7, copied below, sets forth the minimum required FEC and interleaving capabilities for the ATU-C transmitter (and the ATU-R receiver), and Table 8-3, also copied below, sets forth the corresponding minimum FEC and interleaving requirements for the ATU-R transmitter (and ATU-C receiver).

Table 7-7/G.992.1 – Minimum FEC coding capabilities for ATU-C

Parameter	Fast buffer	Interleaved buffer
Parity bytes per R-S codeword	$R_F = 0, 2, 4, 6, 8, 10, 12, 14, 16$ (Note 1)	$R_I = 0, 2, 4, 6, 8, 10, 12, 14, 16$ (Notes 1 and 2)
DMT symbols per R-S codeword	$S = 1$	$S = 1, 2, 4, 8, 16$
Interleave depth	Not applicable	$D = 1, 2, 4, 8, 16, 32, 64$
NOTE 1 – R_F can be > 0 only if $K_F > 0$, and R_I can be > 0 only if $K_I > 0$. NOTE 2 – R_I shall be an integer multiple of S.		

The ATU-C shall also support upstream transmission with at least any combination of the FEC coding capabilities shown in Table 8-3.

Table 8-3/G.992.1 – Minimum FEC coding capabilities for ATU-R

Parameter	Fast buffer	Interleaved buffer
Parity bytes per RS codeword	$R_F = 0, 2, 4, 6, 8, 10, 12, 14, 16$ (Note 1)	$R_I = 0, 2, 4, 6, 8, 10, 12, 14, 16$ (Notes 1 and 2)
DMT symbols per RS codeword	$S = 1$	$S = 1, 2, 4, 8, 16$
Interleave depth	not applicable	$D = 1, 2, 4, 8$
NOTE 1 – R_F can be > 0 only if $K_F > 0$ and R_I can be > 0 only if $K_I > 0$. NOTE 2 – R_I shall be an integer multiple of S.		

The ATU-R shall also support downstream transmission with at least any combination of the FEC coding capabilities shown in Table 7-7.

124. G.992.1 specifies that both the ATU-C and ATU-R must support downstream transmission with any combination of the parameter values shown in Table 7-7, and they must support upstream transmission with any combination of the parameter values shown in Table 8-3. *See, e.g.*, *id.* at §§ 7.6, 8.6.

125. G.992.1 and T1.413 Issue 2 use the same superframe definition as T1.413 Issue 1, *i.e.*, 68 data frames followed by a synchronization symbol, which is identical to the synchronization symbol of T1.413 Issue 1. *See* G.992.1, §§ 3.31, 7.4.1.1, 8.4.1.1. Figure 7-5 of G.992.1, copied below, illustrates superframe structure used by both the ATU-C and ATU-R. *Id.*

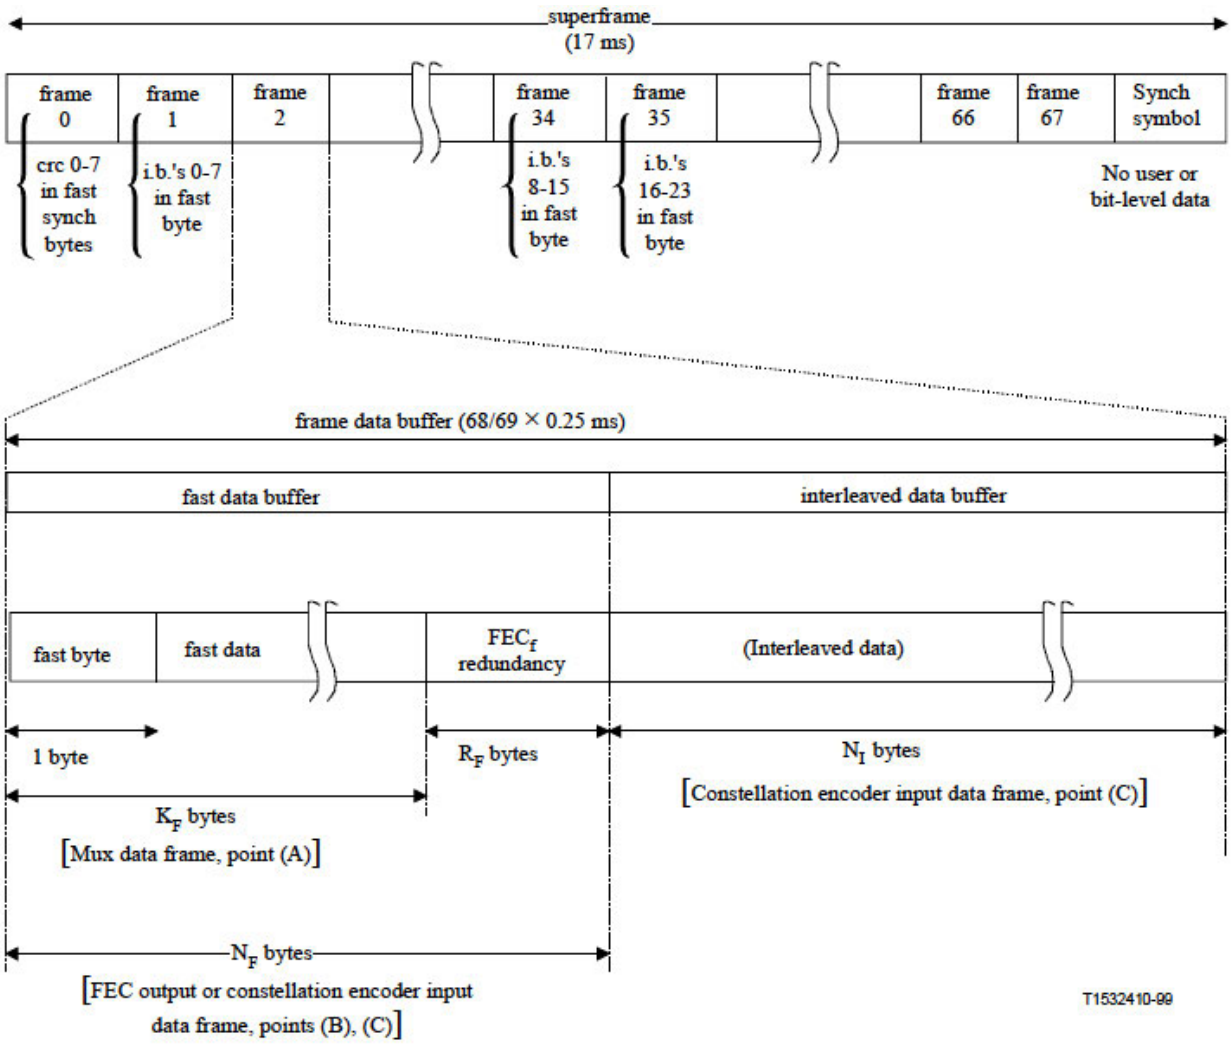

Figure 7-5/G.992.1 – ADSL superframe structure – ATU-C transmitter

126. G.992.1 also defines an "inverse synchronization symbol" that is used in Annex C, which is the annex that specifies aspects of ADSL transceivers deployed in the same binder as TCM-ISDN (e.g., in Japan). G.992.1, §§ C.1, C.4.3.2. The inverse synchronization symbol is used to mark the boundaries between hyperframes, where each hyperframe is five consecutive superframes. *Id.* The inverse synchronization symbol is "a tone-by-tone 180-degree phase reversal of the synchronization symbol (see C.4.7.1) except for the pilot tone." *Id.*; *see also id.* at §§ C.4.7.1, C.5.5.1.

127. Like T1.413 Issue 1, to enable the transmission of both delay-sensitive and delay-tolerant data over the same subscriber line, T1.413 Issue 2 and G.992.1 support dual-latency configurations in which each data frame contains data from two data buffers, as shown in Figure 7-5 above. Delay-sensitive data is put into the fast data buffer for transmission, and delay-tolerant data is put into the interleaved data buffer. *Id.* at § 7.4.1.2. But unlike T1.413 Issue 1, which always defines and uses both a fast buffer and an interleaved buffer, G.992.1 and T1.413 Issue 2 provide for the possibility that the user data requires only one latency path, whether interleaved or non-interleaved. In this case, the amount of overhead can be reduced so that more of the total available bit rate of the connection is available for user data.

128. In the reduced overhead mode introduced in both T1.413 Issue 2 ADSL and G.992.1, the transmitter assigns data only to the fast data buffer or to the interleaved data buffer. When only the fast buffer is used, aoc messages are transmitted in the fast path, and therefore they are not protected by interleaving as they always are in T1.413 Issue 1.

129. G.992.1 and T1.413 Issue 2 accommodate two ways that the two transceivers can transmit simultaneously over the subscriber line: frequency-division multiplexing (FDM) and echo canceling. *Id.* at § 10.1.2 ("Manufacturers may choose to implement this Recommendation using either frequency-division multiplexing (FDM) or echo cancelling (overlapped spectrum) to separate upstream and downstream signals.").

130. Figure 7-1 of G.992.1, copied below, shows the logical channels defined for the ATU-C, and Figure 8-1 of G.992.1, also copied below, shows the logical channels defined for the ATU-R. *Id.* at §§ 7.1.1, 8.1.1.

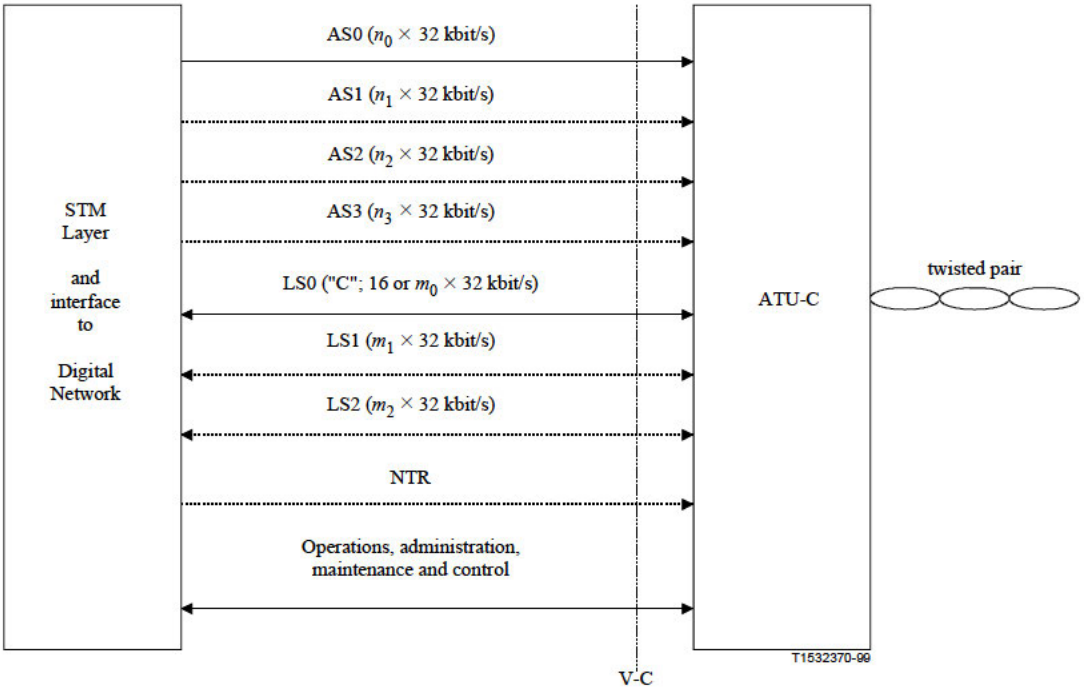

NOTE – Optional bearer channels (both duplex and simplex) and features are shown with dotted lines.

Figure 7-1/G.992.1 – ATU-C functional interfaces for STM transport at the V-C reference point

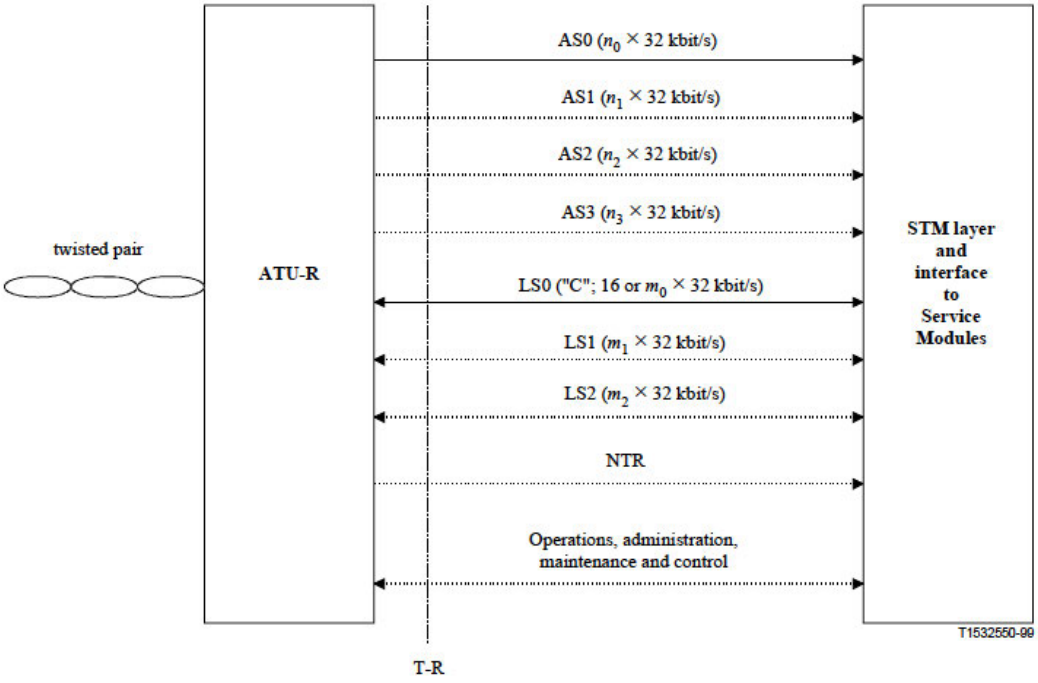

NOTE – Optional bearer channels (both duplex and simplex) and features are shown with dotted lines.

Figure 8-1/G.992.1 – ATU-R functional interfaces for STM transport at the T-R reference point

As Figures 7-1 and 8-1 show, the ATU-C and ATU-R transmit and/or receive up to nine logical channels, all of which are communicated via a single twisted-pair line. To enable the different logical channels to share the same physical line, G.992.1 specifies rules the ATU-R and ATU-C follow to prepare data from all of the bearer channels in use for transmission over the twisted pair. G.992.1 organizes transmitted and received data into data frames. *Id.* at § 1. The ADSL transmitter transmits all of the bits in a data frame in the same DMT symbol. *Id.* at §§ 3.14, 7.4.1.1.

131. G.992.1 specifies initialization procedures that are "required in order for a physically connected ATU-R and ATU-C pair to establish a communications link." *Id.* at § 10.1.1. During initialization, "[i]n order to maximize the throughput and reliability of this link, ADSL transceivers shall determine certain relevant attributes of the connecting channel and establish transmission and processing characteristics suitable to that channel." *Id.* Each transceiver's receiver determines the relevant attributes of the channel through defined transceiver training and channel analysis procedures." *Id.* The initialization procedure also allows the transceivers to establish "[c]ertain processing and transmission characteristics." *Id.* An exchange portion of initialization allows each receiver to communicate to the corresponding far-end transmitter certain transmission settings that it expects to see after the initialization procedure ends and the transceivers begin steady-state communication during Showtime. *Id.* For example, the ATU-C sends to the ATU-R four options for transport configuration for both upstream and downstream, each option including proposed values for the number of FEC redundancy bytes (R), the interleave depth (D), and the number of DMT symbols per FEC codeword (S). *Id.* at § 10.8.3.

a. On-Line Reconfiguration: Bit Swapping

132. Like T1.413 Issue 1, G.992.1 and T1.413 Issue 2 define bit swapping. The protocol is, in all respects relevant to the '835 and '162 patents, identical to the protocol used in T1.413 Issue 1. *See, e.g.*, G.992.1, § 11.2. In the reduced overhead mode, however, when only the fast path is used, the aoc channel may be less reliable because it is not protected by interleaving. G.992.1, § 7.4.3.2.

b. On-Line Reconfiguration: Dynamic Rate Adaptation (DRA)

133. Appendix II of G.992.1, entitled "Dynamic (on-line) Rate Adaptation," describes a mechanism called dynamic rate adaptation (DRA), which "allows reconfiguration of the modem during Showtime" without "a lengthy restart to reconfigure the modem." G.992.1, § II.1.[6] G.992.1 explains that "[t]he purpose of this DRA mechanism is not to provide 'on-the-fly' Rate Adaptation, where the modem configuration would change continuously, tracking the slightest variation of the line conditions without affecting the user-traffic, but rather to allow for occasional changes, which would involve service interruption of the order of tens of milliseconds." *Id.*

134. The described DRA "is a mechanism that during ShowTime, without the need to restart: Allows rate modifications (up and downgrades) for both US and DS." G.992.1, § II.1.1. G.992.1 notes that "[r]ate modification implies more than just bit-rate but also FEC and Interleaving settings." *Id.* A reconfiguration protocol is defined that runs without interfering with user traffic. *Id.* The DRA mechanism does not result in an error-free reconfiguration. *Id.* Instead, "[d]uring the transition period, user-data may be lost during tens of millisecond for both communication directions." *Id.*

[6] Annex K of T1.413 Issue 2 also defines dynamic rate adaptation.

135. The protocol "expands the AOC [(ADSL overhead channel)] message set" and "define[s] new DRA-AOC messages" that allow the ATU-C to "propose to the ATU-R a new rate configuration." *Id.* at § II.2.1. Other aoc messages are used by the transceivers to "exchange configuration information – if the proposal is accepted by ATU-R." *Id.* Finally, the DRA-AOC messages allow the transceivers to "initiate and synchronize a swap to the new rate configuration." *Id.*

136. Two messages allow the ATU-C to propose a new configuration to the ATU-R. G.992.1, § II.4. The DRA_Configuration_Request message sent by the ATU-C includes proposed values for the number of parity bytes per DMT symbol (i.e., number of Reed-Solomon redundancy bytes per DMT symbol) in the downstream fast buffer (R_{fd}), the number of parity bytes per DMT symbol in the downstream interleave buffer (R_{id}), the number of parity bytes per DMT symbol in the upstream fast buffer (R_{fu}), the number of parity bytes per DMT symbol in the upstream interleave buffer (R_{iu}), the number of downstream DMT symbols per FEC codeword (S_d), the number of upstream DMT symbols per FEC codeword (S_u), the downstream interleave depth in codewords (I_d), and the upstream interleave depth in codewords (I_u). *Id.* at § II.4.1. The DRA_Configuration_Reply sent by the ATU-R in response to the DRA_Configuration_Request message indicates whether the ATU-R accepts the proposed new configuration. *Id.* at § II.4.2.

137. After the transceivers have agreed to implement the new configuration, "the swap to the new configuration must be activated and synchronized." G.992.1, § II.6. The ATU-C sends the DRA_Swap_Request message "to inform the ATU-R about when to swap the rate." *Id.* A superframe reference number (SFR) identifies "around which superframe boundary the rate swap will occur." *Id.* G.992.1 states that "[i]f the modems operate with the mandatory S-

values, these SFR-references always coincide with codeword boundaries," which "avoids an explicit Reset of the FEC-mechanism." *Id.* The value of SFR is "zero at the first ShowTime symbol and is then increased by one (modulo 256) at each consecutive superframe." *Id.*

3. ITU-T Recommendation G.992.3 (2003)

138. In 2003, the ITU released ITU-T Recommendation G.992.3, known as "ADSL2" and, while in development, "G.dmt.bis." G.992.3 "describes the second generation of ADSL, based on the first generation ITU-T Rec. G.992.1." G.992.3, p. ii. Accordingly, G.992.3 builds on many of the aspects of G.992.1 to enable connections that support at least 8 Mbit/s downstream and at least 800 kbit/s upstream. *Id.* at § 1. G.992.3 also adds a number of features and capabilities. *Id.*

139. G.992.3 "describes the interface between the telecommunications network and the customer installation in terms of their interaction and electrical characteristics" for ADSL. G.992.3, § 1. G.992.3 specifies the characteristics of and requirements for ADSL2 transmitters. *See, e.g.*, *id.* at p. 1 (scope of G.992.3 Recommendation includes "defin[ing] the line code and the spectral composition of the signals transmitted by both ATU-C and ATU-R" and "specif[ying] the transmit signals at both the ATU-C and ATU-R").

140. G.992.3 specifies the use of Reed-Solomon FEC and convolutional interleaving as two of the functions implemented by the physical-medium-specific transmission convergence (PMS-TC) function. G.992.3, § 7.2 ("the ATU transmit PMS-TC function also provides procedures for: . . . insertion of redundancy for Reed-Solomon-based forward error correction; . . . and interleaving of data frames to spread the effect of impulsive impairments on the U interface. These functions are configured by a number of control parameters described in 7.5 to provide application-appropriate FEC protection, latency, and impulse noise immunity for each frame bearer. The values of the control parameters are set during initialization or

reconfiguration of the ATU. The ATU receive PMS-TC function reverses each of the listed procedures so that the transported information may be recovered.") (emphasis added); *see also id.* at § 7.7.1.4 ("The FEC procedure inserts Reed-Solomon FEC redundancy octets to provide coding gain as an outer coding function to the PMD function. The FEC procedure of latency path function #p shall calculate R_p octets from $M_p \times K_p$ input octets. The octets are appended to the end of the input octets in the structure of FEC Output Data Frame at Reference Point B. When $R_p = 0$, no redundancy octets are appended and the values in the FEC Output Data Frame are identical to the input values. For all other values of R_p, the following encoding procedure shall be used to create the R_p octets: The FEC procedure shall take in M_p scrambled Mux Data Frames comprising message octets, m_0, m_1, … , $m_{Mp \times Kp-2}$, $m_{Mp \times Kp-1}$. The procedure shall produce R_p redundancy octets c_0, c_1, … , c_{Rp-2}, c_{Rp-1}. These two taken together comprise the FEC codeword of size $M_p \times K_p + R_p$ octets. The R_p redundancy octets shall be appended to the message octets to form the FEC Output Data Frame at Reference Point B. At the end of the initialization sequence, the FEC Function always starts with the first of M_p Mux Data Frames."); *id.* at § 7.7.1.5 ("To spread the Reed-Solomon codeword and therefore reduce the probability of failure of the FEC in the presence of impulse noise, the FEC Output Data Frames shall be convolutionally interleaved. The interleaver creates the Interleaved FEC Output Data Frames at Reference point C, at the output of the latency path function. This procedure is followed by the frame multiplexing procedure. Convolutional interleaving is defined by the rule (using the currently defined values of the framing control parameters D_p and the derived parameter NFEC.p): Each of the NFEC.p octets B_0, B_1, … , $B_{NFEC.(p-1)}$ in an FEC Output Data Frame is delayed by an amount that varies linearly with the octet index. More precisely, octet B_i (with index i) is delayed by $(D_p - 1) \times i$ octets, where D_p is the interleaver depth.").

141. Figure 7-6 of G.992.3, copied below, illustrates the PMS-TC functions of an ATU (either ATU-C or ATU-R), including the FEC and interleaving functions.

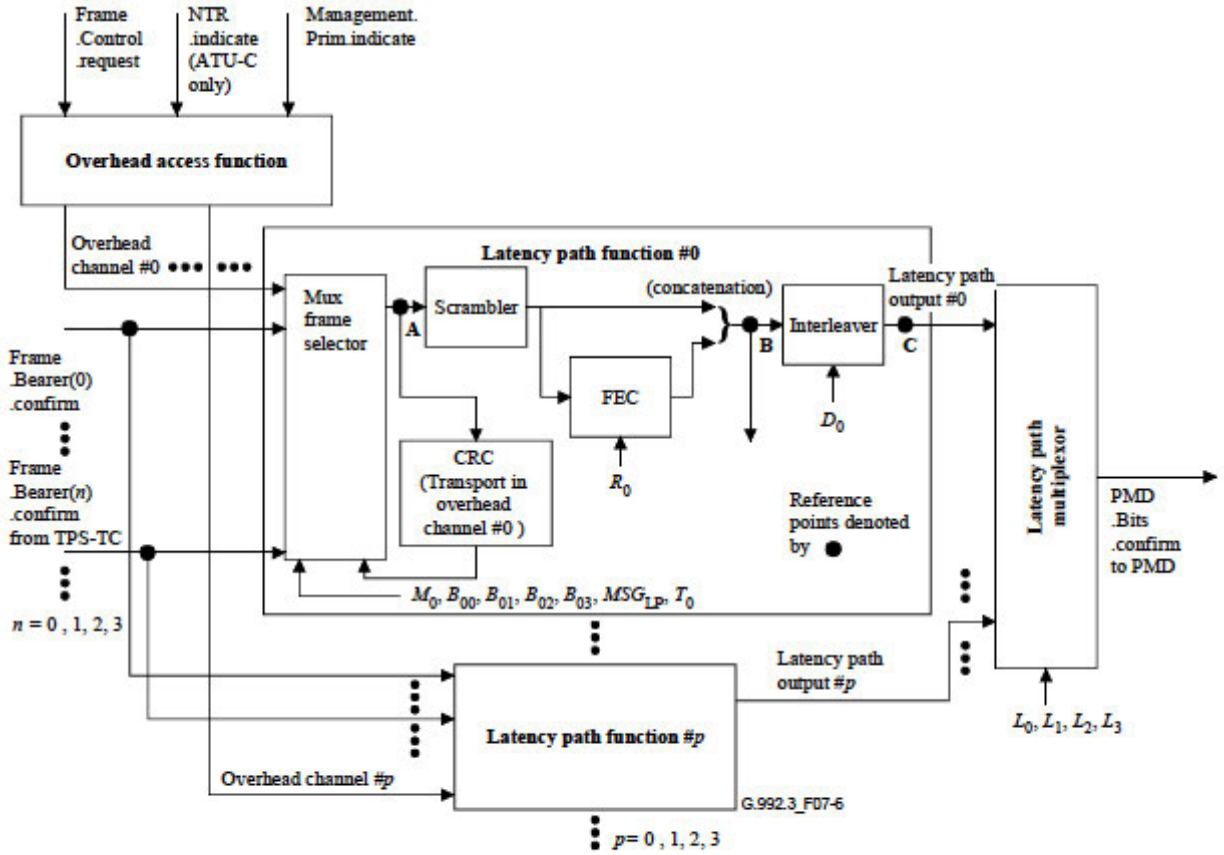

Figure 7-6/G.992.3 – Block diagram of transmit PMS-TC function

142. Table 7-5, copied below, describes the data at the Reference Points A, B, and C.

Table 7-5/G.992.3 – PMS-TC function internal reference points

Reference point	Definition
A: Mux Data Frame	The data within a latency path function after the sync octet has been added.
B: FEC Data Frame	The data within a latency path function after the output of the FEC redundancy octets are merged with scrambled data.
C: Interleaved FEC Data Frame	The data and redundancy octets that have been interleaved. This is the output signal of a latency path function.

143. A set of control parameters specified in G.992.3 controls the configuration of the PMS-TC function, including FEC and interleaving. G.992.3, § 7.5. The control parameters

include R_p, which is the number of Reed-Solomon (RS) redundancy octets per codeword (and per FEC Data Frame) in latency path function #p, and D_p, which is the interleaving depth in the latency path function #p. *Id.* at § 7.5.

144. G.992.3 specifies valid values of each of the control parameters, including R_p and D_p.. The valid values of R_p are 0, 2, 4, 6, 8, 10, 12, 14, and 16, and the valid values of D_p are 1, 2, 4, 8, 16, 32, and 64. G.992.3, § 7.6.2, Table 7-8. If $R_p = 0$ (FEC is off) then $D_p = 1$ (interleaving is off). *Id.*

145. G.992.3 specifies mandatory capabilities for the ATU-C and ATU-R for the mandatory latency path 0. G.992.3, § 7.6.3.1. Specifically, ATU-Cs and ATU-Rs must support all valid values of both R_p and D_p for latency path 0 in both the downstream and upstream directions. *Id.* at § 7.6.3.1, Table 7-9, Table 7-10. If the ATU-C and ATU-R support the optional latency paths, then during initialization the ATU-C and ATU-R identify the maximum values of R_p and D_p they can support on these optional latency paths, and all valid values of R_p and D_p up to and including these maximum values must be supported on the optional latency paths. *Id.* at § 7.6.3.2, Table 7-11, Table 7-12.

146. G.992.3 also defines an overhead channel. G.992.3, § 7.8.2.1. The overhead channel includes a message based portion that transports messages used to coordinate on-line reconfiguration (OLR). *Id.* at §§ 7.8.2.1, 7.8.2.3.

147. G.992.3 characterizes OLR as "a powerful feature of this Recommendation" that allows the ATU-C and ATU-R to "autonomously maintain operation within limits set by control parameters during times when line or environment conditions are slowly changing." G.992.3, § 10.2. OLR also allows the ATU-C and ATU-R to optimize their settings following

initialization. *Id.* Higher-layer data, management, and control functions can make use of OLR. *Id.*

148. Table 9-2, copied below, describes the overhead messages used for OLR, which are designated as being of the highest priority.

Table 9-2/G.992.3 – Highest priority overhead messages

Message and designator	Direction	Command content	Response content
On-line Reconfiguration (OLR) Command $0000\ 0001_b$	From a receiver to the transmitter	New configuration including all necessary PMS-TC and PMD control values.	Followed by either a line signal corresponding to the PMD.Synchflag primitive (not a OLR command) or an OLR command for defer or reject.

149. The OLR messages describe the requested changes to the upstream or downstream TPS-TC, PMS-TC or PMD functions. G.992.3, § 10.2.2. As indicated by Table 9-2, an ATU that receives an OLR command message generates either the "PMD.Synchflag primitive" or another OLR command to defer or reject the command. *Id.* at § 9.4.1, Table 9-2. The PMD.Synchflag primitive causes the ATU to send the Synchflag "as a time marker for when the on-line reconfiguration takes effect." *Id.* at § 10.2.2; *see also id.* at § 10.2.2.1 ("The transmit PMD function transmits the PMD.Synchflag primitive on the line as defined in 8.7, as a time marker for the instant where the reconfiguration will take place.").

150. G.992.3 defines three "types" of OLR messages. G.992.3, § 9.4.1.1. Bit swapping is a mandatory feature that allows the ATU-C and ATU-R to reconfigure the upstream and downstream bits and fine gain (b_i, g_i) parameters without changing any PMS-TC control parameters. *Id.* at § 10.2.1. Bit swapping is coordinated using Type 1 OLR messages. *Id.*

151. Dynamic rate repartitioning (DRR) is an optional feature that, if implemented, can be used to reconfigure the data rate allocation amongst multiple latency paths and can also include modifications to the bits and fine gain (bi, g_i) parameters, reallocating bits among the

subcarriers. G.992.3, § 10.2.1. DRR can change the number of octets from frame bearer #n per Mux Data Frame on latency path #p, i.e., in $B_{p,n}$. *Id.* DRR is coordinated using Type 2 OLR messages. *Id.*

152. Seamless rate adaptation (SRA) is an optional feature that, if implemented, can be used to reconfigure the total data rate by modifying the frame multiplexor control parameters (L_p) and the bits and fine gains (b_i, g_i) parameters. G.992.3, § 10.2.1. SRA is coordinated using Type 3 OLR messages. *Id.*

153. Either the ATU-C or the ATU-R may initiate an OLR. *Id.* The initiating ATU requests changes to the other ATU's transmitter. G.992.3, § 9.4.1.1. In response, the ATU that receives an OLR request may transmit the synchflag (i.e., the "PMD.Synchflag primitive") to acknowledge the requesting ATU's request, or it may defer or reject the request. *Id.* If the response is the PMD.Synchflag primitive, the responding "ATU shall reconfigure the effected PMD, PMS-TC, and TPS-TC functions as described in the reconfiguration clauses describing those functions." *Id.* Likewise, "[u]pon receipt of a line signal corresponding to the PMD.Synchflag primitive, the initiator shall reconfigure the effected [*sic*] PMD, PMS-TC, and TPS-TC functions as described in the reconfiguration clauses describing those functions." *Id.*

154. The Synchflag used in the OLR protocol is transmitted in the last frame of a superframe. As in T1.413 Issue 1, T1.413 Issue 2, and G.992.1, a superframe is a collection of 69 consecutive frames. G.992.3, § 8.4. As shown in Figure 8-6 of G.992.3, copied below, each superframe is composed of 68 data frames, numbered 0 through 67, followed by a synchronization symbol (sometimes referred to as the sync frame). *Id.* Each data frame carries data. *Id.* The synchronization symbol does not carry any data. *Id.* Instead, it is a special symbol, known to the receiver, that helps the receiver identify the superframe boundaries (and thus the

frame boundaries). *Id.* The synchronization symbol is also used to transport the Synchflag used in OLR.

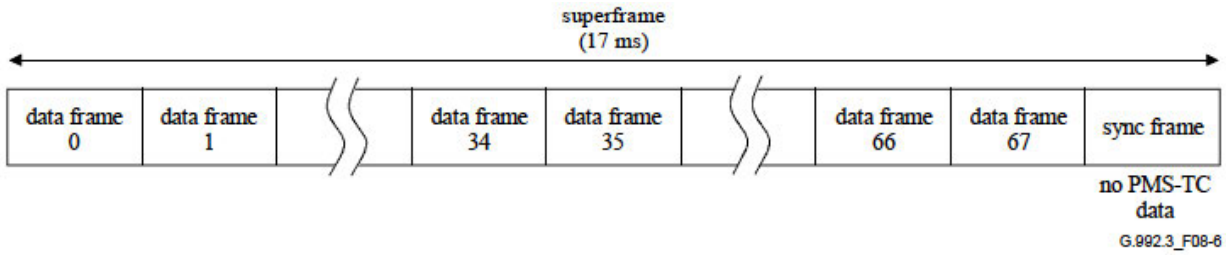

Figure 8-6/G.992.3 – ADSL superframe structure – ATU-C transmitter

155. The synchronization symbol is either a "SS-REVERB" symbol or a "SS-SEGUE" symbol. *Id.* at § 8.7. The difference between a SS-REVERB symbol and a SS-SEGUE symbol is that a SS-SEGUE symbol is a subcarrier-by-subcarrier 180 degree phase reversal of a SS-REVERB symbol. *Id.* In other words, a SS-REVERB symbol modulates the bit pattern generated by the REVERB pseudo-random generator, whereas the SS-SEGUE symbol modulates the bitwise-inverted REVERB bit pattern (i.e., all values of 0 are replaced by 1s, and all values of 1 are replaced by zeros). *Id.*

156. The Synchflag is transmitted by changing the content of the synchronization symbol from its most-recent content. In other words, if the last-transmitted synchronization symbol was a SS-REVERB symbol, then the Synchflag is transmitted by sending a SS-SEGUE symbol, and vice versa. *See, e.g.*, G.992.3, § 8.7.3 ("Every time the transmit PMD function receives a PMD.Synchflag.request primitive (related to on-line reconfiguration during the L0 [Showtime] state) from the transmit PMS-TC layer, the phase of the first next inserted synchronization symbol shall be inverted, and remain inverted until the next PMD.Synchflag.request primitive is to be carried. At the start of Showtime, the first synchronization symbol transmitted shall be an SS-REVERB symbol."). The SS-REVERB

symbol is an inverted sync symbol relative to the SS-SEGUE symbol, and the SS-SEGUE symbol is an inverted sync symbol relative to the SS-REVERB symbol.

157. In the downstream direction, if granted, the requested OLR takes effect starting with the second symbol (symbol index 1) of the next superframe. G.992.3, § 8.16.2. In the upstream direction, if granted, the requested OLR takes effect starting with the fifth symbol (symbol index 4) of the next superframe. *Id.*

158. G.992.3 specifies that when the $B_{p,n}$ values within an existing latency path #p are reconfigured, that reconfiguration occurs "only at boundaries between Interleaved FEC Data Frames." G.992.3, § 7.11.1.1. The timing of the reconfiguration is aligned with "the Interleaved FEC data frame boundary, FEC data frame boundary, Mux Data Frame boundary, and the PMD symbol boundary." *Id.*

E. DSL Standards Contributions

159. The standards described above are the output of the working groups T1E1.4 and SG15/Q4. During development of standards, interested parties submitted written documents, called "contributions," to the working groups. Some of the contents of these contributions were accepted and incorporated into an emerging standard, but many were not. Nevertheless, these contributions reflect what was known in the art and the level of skill in DSL at the times of their submission.

160. In this section, I provide an overview of contributions in several study areas focused on by participants of T1E1.4 and SG15/Q4 before the earliest possible priority date of the '835 and '162 patents.

1. Continuing Goal to Provide Robust, Stable Connections

161. ADSL and VDSL are "always on" services, which means that, absent a serious problem, a connection should not autonomously terminate or reinitialize even if there is no user

data to transmit. *See, e.g.*, T1E1.4/98-279R1, p. 2 ("G.lite and G.dmt have been long since described as 'always on' or 'always connected' services."); BA-085, p. 1 ("The host based ADSL implementations typically consume up to 20% of the CPU cycles for the modem applications. Since the ADSL modems are *always on*, this implies that the CPU will be operating at about 80% of its capability even when there is no payload data transfer."); CI-044, p. 4 ("Since it is suggested that G.dmt and G.lite will operate in an 'always on' or 'always connected' environment, we can assume that the ATU-C will remain powered up at all times (barring extraordinary circumstances) and thus it will be the ATU-R that will determine when a connection will be powered up or powered down."); AB-025R1, p. 1 ("In general we assume that the xTU-C is essentially always on and will have been turned on before the xTU-R. After activation, an xTU-R can remain always on or desire to reenter G.hs for two basic reasons: mobile xTU-R needing to initialize or power being removed or reduced from the xTU-R/PC."); NF-059, p. 2 ("The current belief is that G.dmt and G.lite will operate in an 'always on' or 'always connected' environment and will transport traffic with various delay and throughput requirements.").

162. Because initializing a DMT transceiver from scratch takes a substantial amount of time (typically tens of seconds), it has been a goal in DSL standardization to avoid having to re-initialize a connection at all. When a data interruption cannot be avoided, it is a standing goal to minimize the time required to restore a connection. *See, e.g.*, AB-090, p. 3 ("For some applications, retrains will introduce unacceptable delays."); *id.* at p. 2 ("frequent interruptions due to fast retrains will not be acceptable"); T1E1.4/98-251, p. 2 ("Current standardization efforts suggest a fast retrain of the connection as a means for recovering from the off-hook change. This solution unfortunately causes a loss of service. . . . Disruption of service can be a

problem for DSL. For instance, in internet file transfer, the resulting disruption can lead to a need for retransmission of the data or worse (potential of needing to 'reboot' in some popular PC-based operating systems). Some services may also supply multiple digital POTS services. A disruption of all these digital POTS services everytime [*sic*] the phone rings on the analog POTS could be a problem for potential customers. More examples could be presented, but most customers and service providers view service-disruption as undesirable at least in some situations. Thus, retraining upon off-hook is not a complete solution for splitterless operation."); *id.* at p. 9 ("To avoid a loss of service when line conditions change, we propose a simple seamless solution that uses the existing bit-swapping procedure, hence bypassing the need for fast retrains."); BA-029, p. 27 ("if physical-layer actions of the DSL modem (such as retraining when the margin is low) take place without regard for the state of the higher layer traffic (such as active derived voice calls), this could result in voice calls being dropped. This behavior compromises service feature transparency. The effect of such actions on voice calls in progress MAY be minimized by suitable cross-layer cognizance and communication, such as discussed in 6.2.3. For example, when voice calls are in progress, the DSL modem MAY hold-off retrain until the last possible moment, waiting for an opportunity when no voice calls are in progress. Additionally, if it is known that no voice calls are in progress, retrain MAY be performed to improve margin, even though normal retrain margin threshold has not been crossed. Such selective retraining, which relies on the physical layer management associated with the xDSL modem understanding the nature and state of the payload traffic, could significantly reduce the impact of retraining on derived voice service-quality, whilst subjecting to the interruptions only data traffic with a greater latency/retransmission tolerance. Suitable cross-layer communication MAY therefore also help meet the reliability requirements (section 6.2.11)."); T1E1.4/96-163,

p. 3 ("The alternative is to set the attenuation once, at train-up, and live with systems sometimes retraining when another system comes into use. Such an ungraceful retraining sequence may run a risk of 'pendulum fighting', where one system retraining causes another system to go down, ad infinitum."); AB-033, p. 1 ("There are a number of applications, many of which can be supported over an ATM UBR class of service which will be effected adversely by a excessively long interruptions due to fast retrain events."); *id.* at p. 2 (for video and audio applications, "[i]nterruptions due to fast retrain must be dealt with by buffering as the user can tolerate delay when smoothed out by the buffers. The applications may however be intolerant to data loss or be unable to support significant amounts of retransmission."); T1E1.4/95-057, p. 2 ("If the TDQ and FDQ are retrained after the nullers are turned on, then the DMT system performance improves (Figure 3) and the achieved bit rate is 41.2 Mbps, only about 20% lower than that achieved with no interference. However, retraining results in additional interruption in the data streams and is certainly not desirable in a changing interference environment."); *id.* at p. 3 ("This degradation may be reduced by adjustment of the DMT equalizers to compensate for the new environment, but retraining after every change in the adaptive nulling is not an attractive solution."); D.741, p. 11 ("As soon as the average margin of the system is negative a full re-initialization is required, resulting in a 10 second interruption in service. SRA provides the ability to adapt the data rate seamlessly on-line as the channel changes while always maintaining a positive average margin. This will result in less retrains and generally a more robust connection."); HC-071R1, p. 1 ("If such an error occurs there are a number of recovery techniques that may be employed that involve monitoring slicer error or FEC errors [3], or at worst retraining. However these methods may be undesirable due to the potential impact on modem architecture, complexity, or due to the period of time for which transmission is

interrupted."); T1E1.4/98-279R1, p. 3 ("If for whatever reason, a false transition to data mode occurs, one or more DMT symbols worth of data will either be inserted or deleted from the data stream. Although this erroneous data will be flagged by the CRC as unreliable, the result of inserting or deleting data is loss of ATM cell, FEC codeword and superframe synchronization. Such a loss in synchronization would require a retrain to restore synchronization and may propagate error messages to the higher layers."); T1E1.4/98-306, p. 3 ("Issue: Switchhook transition causes fast retrain (1.5 s of cell loss). Issue: Full initialization causes 10-20 s of cell loss. Issue: Pulse dialing causes service interruption as long as 30s. Issue: When the phone is offhook, the data rate is arbitrarily low (zero is possible). Issue: It was noted that alarms can be raised by these kinds of outages. It was noted that no upstream bandwidth could be assumed during fast retrain."); MA-043, p. 4 ("The suspend method of operations for Qmode has been recognized for its simplicity and minimal impact to current architectures. The primary criticism of the suspend method has been with the required synchronization between the transmitter and receiver and the associated risks of false activation or termination of Qmode by either side. The primary concern is the loss of state information including the FEC codeword, interleaver, and ATM and ADSL scrambler memories due to erroneous Qmode transitions. Such false transitions could be catastrophic to all protocol layers with the worst case requiring fast retrain to reacquire synchronization at the ADSL physical layer."); HC-071R1, p. 1 ("As pointed out in HC-049 a burst of noise might conceivably cause a normal SYNC symbol to be interpreted as an inverted SYNC symbol in which case the receiver could implement the SRA_GO prematurely. HC-049 proposes to check only the symbol on which inversion is expected which eliminates the possibility of a premature implementation of the SRA_GO. However there is still a small possibility that the actual SRA_GO is corrupted by noise such that it appears as a

positive SYNC symbol. If this happens it is presumed that the transmitter would implement the related SRA changes, but the receiver would not! If such an error occurs there are a number of recovery techniques that may be employed that involve monitoring slicer error or FEC errors [3], or at worst retraining. However these methods may be undesirable due to the potential impact on modem architecture, complexity, or due to the period of time for which transmission is interrupted."); SC-060, p. 6 ("Situations where the ATU requesting the OLR either false detects or misses an OLR Acknowledgement are catastrophic. In these events, one and only one ATU will initiate a reconfiguration and this will ultimately lead to a retrain.").

163. The ADSL standards also state that retraining is undesirable and is to be avoided. For example, the synchronization symbol, discussed above, was included in the ADSL standards to "permit[] recovery of the frame boundary after micro-interruptions that might otherwise force retraining."). T1.413 Issue 1, § 6.9.3, § 7.9.3; *see also* T1.413 Issue 2, § 6.11.3 (same), § 7.11.3 (same); G.992.1, § 8.8.1.2 ("a gross timing error that is an integer multiple of this number of samples could still persist after a micro-interruption (e.g., a temporary short-circuit, open circuit or severe line hit); correction of such timing errors is made possible by the use of the synchronization symbol defined in 8.7.").

164. Thus, an ADSL connection can be ongoing for days, weeks, or even months at a time. The "always on" nature of ADSL and the need for connection stability presents some challenges that have been addressed by participants in T1E1.4 and SG15/Q4. These challenges include a need to adapt to changing conditions as the connection ages, and a need to reduce the power consumption of "always on" transceivers.

2. Power Conservation

165. Because ADSL is an "always on" service, during the development of G.992.1 and G.992.2, concerns were raised about transceivers' power consumption. *See, e.g.*,

T1E1.4/98-279R1, p. 2 ("G.lite and G.dmt have been long since described as 'always on' or 'always connected' services. As CPE and CO equipment vendors begin to anticipate the mass deployment of G.lite and G.dmt equipment, the power requirements of such 'always on' equipment has become a concern. CPE equipment vendors are required to meet various 'Green PC' specifications in order to integrate their equipment into future consumer PCs. CO equipment vendors as well as service providers are concerned about the power and associated air conditioning and real estate costs required to deploy xDSL equipment for large numbers of customers. This is further exasperated in remote terminals where power, air conditioning and real estate are extremely costly. If power cannot be controlled, the likelihood of wide spread xDSL deployment may diminish significantly."); MA-043, p. 2 ("Power management of ADSL lines has been an active work item in the ITU-T and has been discussed in ANSI T1E1.4 as well as other committees."); *id.* at p. 3 ("The motivation for quiescent mode is to reduce overall power consumption in a manner transparent to the higher layers connected to the ADSL physical link.").

166. It was recognized that there are times when there is no user data to transmit over an established connection, and that DSL transceivers could reduce power consumption during such times. *See, e.g.*, AB-080, p. 2 ("It is anticipated that considerable additional savings can be achieved by the physical layer recognizing and exploiting periods during which there is no data to transmit. The Quiescent Mode proposes a method of operating in a low processing and low AFE power state when there is no data available for transmission.").

167. Thus, at least by 1998, SG15/Q4 considered adding an "idle state" to ADSL to save power and/or other resources when there is no user data to be transmitted. *See, e.g.*, CI-051, p. 1 ("This paper proposes a very simple, yet robust method for resource conservation

during idle periods in data transmission. During idle mode operation, the transmitter performs little or no processing while the receiver performs reduced complexity time-domain processing."); NF-059, p. 1 ("At the Chicago ITU meeting, it was agreed as a goal to include an 'idle' or 'sleep' mode for G.lite. The motivations expressed for such a mode emphasized resource management and include power savings, modem CPU resource sharing and reduced crosstalk interference."). This proposed "idle state" or "idle mode" eventually morphed into what was referred to as "quiescent mode," "Qmode," and eventually the "L2 state" that was included in G.992.3. *See, e.g.*, G.992.3, p. 3; *id.* at §§ 6.8.1, 7.12.

168. Everyone who proposed to allow the transceivers to transition to and from a lower-power state recognized that the transmitter and receiver would need to coordinate their transitions in order not to adversely affect the user's experience and/or impose substantial additional complexity on the transceivers. *See, e.g.*, CI-051, p. 2 ("Negotiation between the transmitter and receiver is required in order to signal when to enter and exit idle mode. Although expedient entry into idle mode may be desired, the latency of such a transition has no effect to the user. In fact, legacy modems not supporting idle mode would simply never enter idle mode (infinite latency). The same is not true when exiting the idle mode. In this case, any introduced latency greater than one DMT symbol period would be directly applied to the user data and would force buffering at the transmitter. 3Com notes that due to the full duplex requirements of the indicator bits, both the upstream and downstream directions must be idled simultaneously and on superframe boundaries."); T1E1.4/98-279R1, p. 2 ("ATM cell synchronization must be maintained during Qmode operation and transitions to and from Qmode. A loss of cell synchronization at any time would result in dropped cells, which is unacceptable for transparent operation."); *id.* ("ATM cell synchronization must be maintained during Qmode operation and

transitions to and from Qmode. A loss of cell synchronization at any time would result in dropped cells, which is unacceptable for transparent operation."); *id.* at p. 3 ("If for whatever reason, a false transition to data mode occurs, one or more DMT symbols worth of data will either be inserted or deleted from the data stream. Although this erroneous data will be flagged by the CRC as unreliable, the result of inserting or deleting data is loss of ATM cell, FEC codeword and superframe synchronization. Such a loss in synchronization would require a retrain to restore synchronization and may propagate error messages to the higher layers. The only recourse is to provide a robust transition technique that has a very low probability of false transitions (addressed in Section 4)."); MA-043, p. 2 (Qmode requirements include "Qmode operation may not increase the link failure rate."); *id.* (Qmode requirements include "No bit errors nor spurious alarms are caused by Qmode."); *id.* ("Loss of synchronization at the ATM layer due to a Qmode event could cause severe data loss."); *id.* at p. 3 ("Qmode activity at the PHY layer should not degrade data rate or BER performance for active data. Thus, the data rate after Qmode should be no different than if Qmode had not occurred. This requires periodic channel probing during Qmode to track and compensate for changes in the channel as well as adjusting frequency and time domain equalizers in the receiver."); *id.* at p. 4 ("Failure to correctly set the receiver thresholds would result in lost data (if a real DMT symbol is falsely identified as noise and therefore discarded) or incorrectly inserted data (if a Qmode symbol is falsely identified as return to data mode). Either false receiver condition could lead to a loss of cell delineation and subsequent catastrophic data loss.").

169. Accordingly, several approaches to coordinating transceivers' transitions into and out of lower-power modes were proposed and discussed at standards meetings starting in 1998. To reduce complexity and take advantage of existing elements of ADSL transceivers,

many of these proposals suggested transmitting an inverted synchronization symbol to signal transitions between a lower-power state and the full-power state. *See, e.g.*, CI-051, p. 2 ("This proposal suggests that during idle periods in data, a single DMT symbol shall be transmitted repetitively. The DMT symbol to be transmitted would be selected such that its spectral properties match those of the data bearing DMT symbols. A 180° phase shift of this DMT symbol would signal the end of the idle period. Therefore, the transmitter would calculate the idle DMT symbol once at the start of the idle period and then simply repeat this symbol until the end of the idle period in which case it would invert the final idle symbol (with the exception of the pilot tone). At the receiver, an idle DMT symbol detector is implemented (this can be performed in the time-domain) along with a phase detector. Upon detection of a 180° phase shift of the idle symbol, the receiver would exit the idle mode. 3Com proposes that the superframe synchronization symbol be used as the idle symbol. . . . The selection of the superframe synchronization symbol implies that no additional complexity is required as both the symbol generator and 180° phase shifter are currently specified in the standard. Furthermore, this symbol is currently used during initialization and can therefore be used to refine any adaptive elements in the receiver without any additional complexity."); *id.* at p. 3 ("For exiting idle mode, 3Com proposes that the final idle mode symbol be shifted 180° with respect to the other idle mode symbols (with the exception of the pilot carrier which may be required for timing recovery). This can be initiated by either the ATU-C or ATU-R. Upon detection, the receiver assumes the next symbol is the first symbol of the first superframe following the idle period interruption."); NF-059, p. 1 (requesting agreement "That the idle mode shall use a 180 degree phase shifted idle symbol to indicate the end of the idle state."); *id.* at pp. 5-6 ("When in the idle state, the transmitter shall repetitively transmit a single DMT

symbol. The DMT symbol to be transmitted is selected such that its spectral properties match those of the DMT symbols transmitted in the data state (note that this could be either the L0 data state or the optional L1 data state as proposed for power management). A 180° phase shift of this DMT symbol would signal the end of the idle period (if a pilot tone is supported for loop-timing, the pilot carrier shall not be phase shifted). This contribution suggests that the idle symbol be identical to the superframe synchronization symbol for these reasons: The superframe synchronization symbol is already defined as part of the standard. By re-using this symbol, the complexity is kept to a minimum. The superframe synchronization symbol is a wide-band signal with the same spectral properties as the data symbols and thus meets the requirements for the idle signal. The superframe synchronization symbol is the last symbol to be transmitted during the data state and therefore does not need to be recalculated when entering the idle state. The superframe synchronization symbol belongs to the same family of symbols used during initialization and therefore can be used to refine any adaptive elements in the receiver without any additional complexity."); *id.* at p. 6 ("When operating in the idle state, the receiver shall examine each DMT idle symbol and be able to distinguish the idle symbol from the end-of-idle symbol."); WH-077, p. 5 ("The idle symbol is the superframe synch symbol and the 180 degree phase shifted superframe synch symbol is the end-of-idle symbol. The 180 degree phase shifter and phase shift detector are already implemented as part of the REVERB to SEGUE transitions in the initialization routine. Finally, the negotiation process as performed in the AOC, is very akin to the dynamic rate adaptation function already implemented and thus would only require the definition of two new messages. Thus, the only additional complexity would be the logic required to detect idle periods in data transmission, generate and process negotiation messages and to indicate if the idle or end-of-idle symbol should be transmitted.");

AB-028, p. 2 ("The dynamic powersave method needs a wakeup method which can be used at any time during the superframe. It can have a gotosleep method aligned with the superframe boundary. For example, an eoc or aoc message or a SEGUE in frame 69 can indicate gotosleep. A REVERB in any frame can indicate wakeup. REVERB and SEGUE symbols be distinguished from one another at a given position and from a QUIET symbol at any time, with error probability of false and non-detection of e.g., 1E-7. *Conclusion: Low to high data rate switching can be reliably achieved with insertion of SEGUE and/or REVERB symbols. The same method or a message based method may be used for high to low data rate switching.*"); AB-080, p. 2 ("The quiescent mode applies to the L0 state only and, in essence, can be considered as a new L0q state. . . . During operation in the L0q state, only every 69th DMT symbol (the superframe synchronization symbol) is transmitted. It is intended that this 69th symbol be used to monitor the channel for on-hook/off-hook transitions and other line anomalies. For the other 68 DMT symbols, nothing is transmitted in the upstream direction and only the pilot is transmitted in the downstream direction."); *id.* at p. 3 ("At L0q state, the physical layer is suspended in a low MIPS state and the average AFE driver power is reduced. Only every 69th symbol, being the superframe synchronization symbol, is transmitted. For the remaining 68 symbols, nothing is transmitted in the upstream direction, and only the pilot is transmitted in the downstream direction. By maintaining the superframe boundaries at 17 msec intervals, minimizes the changes required. The use of the superframe synchronization symbol allows for the periodic detection of on-hook/off-hook transitions or other channel anomalies as well as provides a means of maintaining state synchronization within the power management state machine. When either ATU detects user data, as determined at the ATM TC layer, or any EOC or AOC messages, it shall generate a wakeup signal to switch the L0q state back to the full rate L0 state.

The delay in returning to the L0 state is negligible and ensures the transition back to L0 is transparent to the higher layers. . . . The wakeup from L0q to L0 is performed via a phase reversed superframe synchronization symbol [1]. This symbol may be transmitted during any idle DMT symbol period (not just at the 69th symbol). This ensures negligible latency in the state transition. Immediately following this 'wakeup' symbol, the first data symbol (i.e. DMT symbol 0) of the next superframe is transmitted."); *id.* at p. 5 ("Once the ATUs have successfully completed the EOC negotiation process, both ATUs shall transition to Quiescent Mode on the superframe boundary specified in the Quiescent_Request message. . . . The transition from the Quiescent Mode to full data rate is commanded using the QMCMD symbol. Upon entry into the Quiescent Mode, the physical layer and ATM TC sub-layer at both ATUs shall retain their current full data rate state. Upon exit from Quiescent Mode, the ATUs shall resume from this previous state as if the ATUs had never entered the Quiescent Mode. The superframe reference number shall not be reset, the Reed Solomon code word synchronization shall be maintained and the TC sub-layer synchronization shall be maintained. Table A.4 defines the set of DMT symbols that may be transmitted during Quiescent Mode operation. . . . Every 69th DMT symbol is the QMSYNC symbol. The position of the QMSYNC symbol shall align to the position of the superframe synchronization symbol in the full rate data mode. Between QMSYNC symbols, 68 QMIDLE symbols shall be transmitted. Either ATU may command a return from Quiescent Mode using the QMCMD symbol. This symbol may be transmitted during any symbol period and the symbols following the QMCMD are the first data symbols of the next superframe."); *id.* (Table A.4 indicates that C-QMSYNC is "Identical to the ATU-C superframe synchronization symbol," C-QMCMD is "Identical to the C-QMSYNC symbol with all carriers except the pilot (sub-carrier 64) phase-shifted by 180°," R-QMSYNC

is "Identical to the ATU-R superframe synchronization symbol," and R-QMCMD is "Identical to the R-QMSYNC symbol with all carriers phase-shifted by 180°") (emphases added).

170. Thus, the use of an inverted synchronization symbol in a protocol to coordinate transmitter and receiver transitions from a lower-power state to a higher-power state was known and had been proposed in SG15/Q4 nearly six years before the earliest possible priority date of the '835 and '162 patents.

171. In January of 1999, Alcatel submitted a contribution proposing a quiescent mode. PO-038R1. Among the benefits of the proposal were that "[t]he return to normal (non-quiet) operation can be done at any time in a seamless way, transparently to full data load, without signalling," "[f]or interleaved data transport, proper flushing of the interleaver memory is ensured, without any loss of data," "[a]ll data (both ATM and the framing overhead, such as AOC, EOC, CRC, Indicator bits) can go from an idle state to activity transparently," and "[u]nder impulsive noise the erroneous return from Qmode to L0 [non-quiet operation] will have a better noise immunity, than in the permanent L0 state." *Id*. at p. 3.

172. Alcatel noted that the entry into Qmode could be made at a convenient point, namely, at the beginning of a FEC codeword: "the Generic TC layer at the transmitter could go to Qmode at the beginning of every new FEC codeword." PO-038R1, p. 5. But because an exit from Qmode would be triggered by a need to transmit user data, which could occur at any time, such as "in the middle of a codeword (in the case of interleaved data transport with S>1)," the transmitter might need to generate a FEC codeword containing some idle data and some user data. *Id.* at p. 6 ("it is possible that a RS code block contains idle (= zero) data first, send with a Qmode symbols, followed by the start of a non-zero data burst."). For this reason, "reconstruction is tricky!!" *Id.* To avoid codewords including a mix of real, non-zero data (e.g.,

user data) and Qmode symbols (containing no data), Alcatel observed that "[i]f we require that after each transition from idle to active or from active to idle at least one full code word is sent, without puncturing it with Qmode symbols, then the receiver will see this unpunctured RS code word, wich [*sic*] enables a likely correct update of the state of the [self-synchronizing scrambler]." *Id.* A simpler solution, Alcatel noted, would be to delay the transition back to the full-power state until the beginning of the next FEC codeword boundary. *Id.* ("When the TC layer leaves the quiet state one could require that a full RS code word with all idle (=zero) data is completed, and that the next RS code-word, containing the transition from zero to non-zero data is not punctured by Qmode symbols. In other words: to avoid the decoding of partially quiet R/S code-words, one forbids the partial (quiet) transmission of any RS code-word, when it is not containing all data bytes (before scrambling) at zero. This adds extra delay, but is [*sic*] allows the receiver to discard all RS blocks received with a number of missing bytes, sent via a quiet symbol, and to track the state of the descrambler, by predicting only zeros.").

173. In other words, Alcatel recognized that forcing Qmode transitions to occur only at the boundaries of FEC codewords would be desirable because it would simplify the implementation.

3. On-Line Reconfiguration

174. It has been recognized since the earliest days of ADSL that as Showtime progresses, aspects of the communication environment can change. For example, the noise on the subscriber line can fluctuate, or it can become substantially worse than it was when Showtime first began. If these changes in the communication environment are significant enough, the receiver will make detection errors, which can cause the connection to fail. Consequently, as described in further detail below, the ADSL standards have always specified

procedures to make changes to the transmitter and receiver to compensate for changes in the communication environment.

175. In addition, the type of traffic (data) transmitted over a connection can change after the initialization procedure has been completed. For example, DSL connections can simultaneously provide for the transfer of traffic for Internet access, video streaming, and voice calls, depending on customer needs. As is well known, the quality-of-service requirements for different types of traffic differ. For example, voice calls are highly sensitive to delay (also referred to as latency) because large delays are noticeable and disconcerting to listeners, but voice calls are somewhat insensitive to data loss because listeners can still understand what has been said as long as enough of the speech arrives intact. In contrast, video streaming is highly sensitive to data loss because the eye notices artifacts in reconstructed video streams, but it is less sensitive to delay because the receiver includes a buffer that can hold a few seconds of video prior to playback. Thus, DSL transceivers are able to provide multiple latency paths having different characteristics (e.g., delay, bit rate, etc.) that allow the transport different types of data simultaneously over the same connection. For example, a telephone call supported digitally over a DSL connection can be carried by a low-latency, low-bit-rate latency path, which might have very little or no FEC and little or no interleaving, whereas streaming of a Netflix video can be carried by a high-bit-rate, high-latency latency path, which would have a more substantial degree of FEC and would also use interleaving to mitigate errors in the decoded video stream.

176. As explained above in the discussions of T1.413 Issue 1, T1.413 Issue 2, and G.992.1, in order to accommodate changes to the communication environment and/or data to be transported without having to perform some kind of re-initialization procedure, DSL

standards specify procedures to make changes to the transmitter and receiver settings after the transceivers have entered Showtime. These procedures are known generally as "on-line reconfiguration (OLR)."

177. As described above, prior to the '835 and '162 patents' earliest possible priority date, the DSL standards specified several types of OLR. T1.413 Issue 1 specifies bit swapping. G.992.1 and T1.413 Issue 2 define bit swapping and dynamic rate adaptation. The 2002 version of G.992.3, which was in existence on the '835 and '162 patents' priority date, defines three types of OLR: bit swapping, dynamic rate repartitioning, and seamless rate adaptation.

178. All of the OLR types require some kind of robust protocol to coordinate changes so that both the transmitter and receiver transition to a new configuration at the correct time. Otherwise, the receiver and transmitter settings become mismatched, which at least causes the receiver to make detection errors and, in the worst case, causes the transceivers to drop the connection and perform the entire initialization sequence again, thereby interrupting data transfer for at least tens of seconds. As explained above, retraining has always been considered to be undesirable in DSL.

a. Bit Swapping

179. The original form of OLR, which continues to be required in all ADSL and VDSL standards, is bit swapping. Bit swapping, which I described above in Sections VII.D.1.d, VII.D.2.a, and VII.D.3, maintains the bit rate of the connection but allows the transmitter and receiver to adjust how many bits are carried by, and how much power is allocated to, each subcarrier to adapt to current conditions on the subscriber line. Bit swapping does not change the data rate of the connection. Furthermore, if a connection supports more than one latency path, bit swapping does not change the bit rate of any path.

180. The ability to conduct bit swapping is mandatory in all of the ADSL standards in order to avoid the transceivers having to retrain because of slow changes to the subcarrier SNRs. *See, e.g.*, T1E1.4/98-261, p. 3 ("Bit-swap should be mandatory in the G.lite Recommendation. Bit-swap has been implemented in ANSI T.413 compliant ADSL systems, and has proved to be valuable to assure the good loop-reach and bit-rate performance. Without bit-swap, the only recourse to detrimental changes in transmission environment is to adapt to a lower bit-rate or retrain. The methods for dynamic rate change are complex and are unlikely to be ready for approval soon. Fast retrains are designed to cope with on/off hook events and are poorly suited for slowly varying line conditions. A full retrain interrupts transmission for about 10 seconds. Without bit-swap, a greater safety margin for Signal-to-Noise-Ratio would likely be necessary to prevent excessive rate-change or retrain events. Thus, the nominal performance would be reduced.").

b. Dynamic Rate Adaptation

181. The DRA mechanism defined in G.992.1 and discussed above served as a springboard for development of dynamic rate repartitioning and seamless rate adaptation techniques, which were later added to G.992.3. Alterations of the DRA protocol were also proposed and discussed for use with other features and applications, such as idle/quiescent mode (discussed above) and to support voice over DSL (VoDSL). *See, e.g.*, FI-133, p. 1 ("In this contribution, we suggest a scheme for seamless rate adaptation, which does not require any changes in the standard ADSL framing, and utilizes the DRA protocol, which was already defined by the standard."); D.817, p. 1 (same); BI-071, p. 1 ("This contribution proposes a mechanism for Dynamic Rate Repartitioning (DRR) for Channelized Voice over DSL (CVoDSL) systems. It is based on the current DRA messages and will be adapted to G.bis framing messages."); HC-038, p. 4 ("While the actual mechanism for dynamic bandwidth

changes is dependent upon the framing of the physical layer and is to be defined, Dynamic Rate Adaptation (DRA, G.992.1 appendix II) can be considered for dynamically altering the data bandwidth partitioning between ADSL data channels."); BA-055R1, p. 3 ("Dynamic Rate Adaptation (DRA, G.992.1 appendix II) provides a mechanism for dynamically altering the data bandwidth partitioning between ADSL data channels."); D.694, p. 5 (same); NF-059, p. 3 ("Negotiation to enter the idle state from the data state is performed via the ADSL overhead channel (AOC). The AOC channel is appropriate for this purpose as the idle mode is very akin to the dynamic rate adaptation (DRA) function and would not further reduce the limited number of available EOC channel opcodes. This contribution proposes that the two commands listed in Table 1 be added to the DRA AOC command set. . . ."); IC-079R1, p. 2 ("BI-071 proposed using a simplified DRA to avoid some of the problems with using DRA to accomplish DRR.").

c. Dynamic Rate Repartitioning

182. Dynamic rate repartitioning (DRR), added as an option to G.992.3, is used to reconfigure how the total data rate of a connection is distributed among multiple latency paths. The total bit rate of the connection remains the same, but the allocation of that total bit rate among different latency paths in use can be changed through DRR. *See, e.g.*, HC-038, p. 4 ("The data bandwidth of an ADSL connection is limited, particularly on long loops and/or in the upstream direction. Multiple 64kbps PCM data channels can use a substantial fraction of the available ADSL data bandwidth. As a result, it is highly desirable to be able to dynamically repartition the data bandwidth between the ATM and STM ADSL data channels.").

183. Texas Instruments was one of the companies that proposed to use the DRA protocol defined in G.992.1 and T1.413 Issue 2 to coordinate DRR. *See, e.g.*, HC-038, p. 4 ("While the actual mechanism for dynamic bandwidth changes is dependent upon the framing of the physical layer and is to be defined, Dynamic Rate Adaptation (DRA, G.992.1 appendix

II) can be considered for dynamically altering the data bandwidth partitioning between ADSL data channels. Such an example reveals many of the key issues associated with dynamic bandwidth changes for VoDSL systems. An information exchange protocol such as the DRA handshake is needed to support dynamic bandwidth changes. However, in contrast with the present DRA protocols, VoDSL systems may benefit significantly from reduced complexity bandwidth modification schemes which reallocate data between ATM and STM channels without modifying the aggregate xDSL data rate or modulation parameters. This simplified bandwidth modification protocol is termed Dynamic Rate Reallocation (DRR).").

d. Seamless Rate Adaptation

184. One form of seamless rate adaptation (SRA) was eventually added to G.992.3 as an option and is used to adjust the total data rate of the connection. Unlike DRR, which merely moves bits from one latency path to another but keeps the total number of bits transmitted per second the same, the optional standardized version of SRA changes the total data rate of the connection. But when SRA was being discussed in SG15/Q4, it was envisioned to have broader use, including to coordinate transitions between full-power and lower-power modes.

185. For example, at the SG15/Q4 meeting held from January 31 through February 4, 2000, a group of companies presented FI-133, entitled "Proposal for seamless rate adaptation based on standard ADSL framing and on the DRA algorithm." The contribution proposed "a scheme for seamless rate adaptation, which does not require any changes in the standard ADSL framing, and utilizes the DRA protocol, which was already defined by the standard." FI-133, p. 1. The proponents of FI-133 proposed to allow the size of the FEC codeword, the number of redundancy bytes per FEC codeword, and the number of DMT symbols per codeword (S) to be changed via SRA: "The DRA algorithm can be used without loss of data, if the parameter N of

the interleaver is decoupled from the parameter N_{FEC} of the FEC. The parameters of the interleaver (D and N) can be determined on startup, and not changed when the rate is changing. This way, the operation of the interleaver does not depend on the rate. The functionality of the interleaver is similar to the definition in the standard. Changing parameters L, N_{FEC}, S and R changes the rate without a loss of data." *Id.* at p. 1. Among the advantages of the proposal was that "[t]he DRA algorithm can be utilized with minor modification." *Id.* at p. 2.

186. At the April 2000 meeting in Geneva, Orckit proposed substantially the same proposal as presented in FI-133. *See, e.g.*, D.817, p. 1 ("In this contribution, we suggest a scheme for seamless rate adaptation, which does not require any changes in the standard ADSL framing, and utilizes the DRA protocol, which was already defined by the standard. This scheme does not require a decoupling of the framing from the physical level. . . ."); *id.* at p. 2 ("The DRA algorithm can be used without loss of data, if the parameter N of the interleaver is decoupled from the parameter N_{FEC} of the FEC. The parameters of the interleaver (D and N) can be determined on startup, and not changed when the rate is changing. This way, the operation of the interleaver does not depend on the rate. The functionality of the interleaver is similar to the definition in the standard. Changing parameters L, N_{FEC}, S and R changes the rate without a loss of data."); *id.* at p. 3 (advantages of proposal include "The DRA algorithm can be utilized with minor modification."); *id.* ("We have described a seamless rate adaptation method that is based on the DRA protocol. We propose that we agree to adopt the seamless rate adaptation scheme described in this document.").

187. At the same meeting, Mr. Tzannes proposed "normal SRA" (NSRA) and "fast SRA" (FSRA) protocols. D.741, pp. 5-8. Both the NSRA and FSRA protocols used "an inverted sync symbol as a flag to signal the receiver" to transition to a different configuration. *See, e.g.*,

id. at pp. 5-7. The contribution proposed to allow the transmitting transceiver to "send[] the inverted sync symbol as a signal to indicate the immediate transition into low power mode," and then the receiver "detects the inverted sync symbol that indicates the immediate transition into low power mode." *Id.* at p. 9. D.741 also proposed to rely on "the inverted sync symbol" to exit low power mode. *See, e.g.*, *id.* at p. 9 ("If the transmitter has not turned off the sync symbol in low power mode the NSRA or FSRA protocols would be used as described above. If the transmitter had turned the sync symbol off while in low power mode, then the 'SRA Go' is sent by the transmitter by turning the sync symbol back on. In this case the receiver would simply try to detect the presence of the sync symbol (with or without inversion) as a flag to synchronize the exit from the low power mode."); *id.* at p. 7 ("In step 7 the transmitter sends the inverted sync symbol as a signal to indicate the immediate exit from low power mode. In step 8 the receiver detects the inverted sync symbol that indicates the immediate exit from low power mode.").

188. At the SG15/Q4 meeting held in August of 2000, with respect to the use of an inverted sync symbol, Alcatel asked, "Should the SRA swap message include indication of at which superframe the SRA swap must take place?" HC-049, p. 2. In response, Alcatel concluded that "[t]he receiver must not monitor all sync symbols but just the particular one requested. Then only that inverted sync symbol is to be monitored. No chance of taking a noise on another one as inverted sync. Matter of robustness. Annex C situation may be better of with this. Then a non-hyperframe boundary can be indicated for inverted sync symbol as swap GO." *Id.*

189. Also in August of 2000, Nortel Networks proposed "an alternative method of using the inverted SYNC symbol to make SRA more robust." HC-071R1, p. 1. Specifically,

Nortel noted that "a burst of noise might conceivably cause a normal SYNC symbol to be interpreted as an inverted SYNC symbol," which could cause the receiver to implement SRA prematurely. *Id.* To make this circumstance less likely and improve robustness, Nortel proposed that the transmitter would "invert[] the current state of the sync symbol" each time a SRA was to be implemented. *Id.* Starting from the beginning of Showtime, "only normal SYNC symbols would be transmitted until the 1st SRA command is implemented, then only inverted SYNC symbols would be transmitted until the 2nd SRA command is implemented, etc." *Id.* This approach was eventually adopted in G.992.3 for the OLR protocol.

4. Modifications to FEC and Interleaver Parameters During Showtime

190. As discussed above, the SRA protocol proposed by Orckit and others in FI-133 and D.817 allowed seamless (error-free) changes to FEC parameters. Other contributions also discussed how to seamlessly change FEC parameters, and they considered whether interleaver parameters could be changed, seamlessly or not.

191. In June of 2000, Catena Networks described what it referred to as "adaptive framing" in the fast path. BA-053, p. 1. Specifically, Catena proposed to allow "the FEC and Mux Data Frame parameters to change" after setting them during initialization. *Id.* The proposal offered the benefit of "maintain[ing] the advantages of using the Fast path while still permitting the advantages of rate adaptation and re-partitioning." *Id.*

192. First, the contribution described how adaptive framing could be accomplished without removing the restriction in G.992.1 that "the FEC and Mux Data Frame parameters . . . are aligned to the DMT symbol." BA-053, pp. 1-2. Catena noted that "[i]f the Line Rate (L) for the Fast path is changed during ShowTime, the FEC and Mux Data frame parameters for that path can also be changed to maintain this alignment." *Id.* at p. 2.

193. Next, Catena considered how the proposal would need to be changed if the FEC codewords were not aligned with DMT symbols (as eventually became the case in G.992.3) and concluded that "[i]n order to seamlessly apply these changes, it must be done on a codeword boundary." *Id.* at 3. Referring to the then-current draft of G.992.3, the contribution noted that in the case that alignment between DMT symbols and FEC codewords were not maintained, "worst case periodicity for achieving DMT symbol – codeword alignment if they are not byte or codeword aligned is 2040 DMT symbols or 0.51 seconds." *Id.*; *see also* D.834, p. 3 ("The boundaries of the Constellation Encoder Input Data Frame (i.e., symbol boundary) and the RS codeword are aligned at least at the beginning of SHOWTIME and further every (8*NFEC) symbols. The alignment occurs therefore at least every 8*255=2040 data symbols or 0.51 seconds.").

194. BA-053 also considered how to coordinate the change to new framing parameters and concluded that "[c]oordination of the framing changes can be handled in the same manner as that used for any rate adaptation or rate repartitioning." BA-053, p. 3. Assuming that FEC codeword boundaries remained aligned with DMT symbol boundaries, "all FEC and Mux Data Frame parameters are also able to be changed on data frame and, thus, superframe boundaries." *Id.* The contribution concluded that for "any exchange in which DMT symbol parameters are communicated between receiver and transmitter, the associated FEC and Mux Data Frame parameters must also be exchanged for the Fast path (B, L, M, and/or S)." *Id.* Finally, Catena proposed that SG15/Q4 should agree that "parameter changes be allowed to the Mux Data Frame and the RS FEC codeword coincident with any changes made to the Line rate for the Fast Path." *Id.*

195. At a subsequent SG15/Q4 meeting held from July 31 to August 4, 2000, Alcatel posed, and proposed answers to, some questions related to SRA features and messages. HC-049, p. 1. For example, Alcatel asked, "Should SRA messages also include Mux Data Frame, RS Codeword and interleaver configuration changes?" *Id.* Answering its own question, Alcatel said:

> Then we have save_new_config message. The swap superframe boundary should then coincide with RS codeword and/or Mux Data Frame boundary. For RS, that is worst case every 4.3 sec (if L even). Only a swap at that point can change the RS config param. With interleaving this will be not be seamless if NFEC is tied to NINT. We still have NINT not tied to NFEC. Interleaving is only disturbed when NINT or D is changed. Not when Mux Data Frame or RS is changed. Config change is also better of with superframe boundary to swap indicated in SRA request message. Then it can be verified upfront whether indicated superframe boundary is ok with configuration change (boundary line-up). GO can then not be sent. Otherwise needs verification afterwards (likely to be too late).

Id.

196. Regarding changing interleaver parameters, Alcatel asked, "Should the SRA itself and the seamlessness of the SRA swap be a mandatory or optional capability?" HC-049, p. 2. Alcatel stated, "Given the agreement on S-decoupling, we could require seamless rate adaptation when possible. Seamless adaptation should be required as long as NINT and D are not changed. In latter cases, it is not possible anyway." *Id.* Alcatel also asked, "Should transition to emergency powering be possible with SRA?" *Id.* In answer to its own question, Alcatel said, "NINT and D changes may be required if rate change is huge eg for emergency powering fallback to 64kbit/s and delay must be kept reasonable. This is not seamless but that is no issue when transiting to emergency power. Transition to emergency powering can still have profile swap mechanism in common with SRA mechanism." *Id.*

F. Summary of State of the Art as of the '835 and '162 Patents' Priority Date

197. This section summarizes aspects of the state of the art in DSL as of the '835 and '162 patents' earliest possible priority date of March 3, 2004.

1. Existing Standards

198. As explained above, as of the earliest possible priority date of the '835 and '162 patents, several ADSL standards had been completed and published, including T1.413 Issue 1, T1.413 Issue 2, G.992.1, G.992.2, and G.992.3. Each of these standards was in force, equipment and chipsets supporting some or all of these standards had been developed and was being sold by numerous companies, and those working in the DSL field would have been familiar with and would have referred to some or all of these standards as a matter of course during his or her work.

2. Disclosures in Standards and Contributions

199. Furthermore, before the earliest possible priority date of the '835 and '162 patents, thousands of contributions had been submitted to the various DSL standardization bodies. These proposals disclosed, among other things:

a. A transmitter sending an inverted synchronization symbol as a flag to coordinate its and the receiver's configuration change

200. At least by 1998, 3Com, later joined by other companies such as Texas Instruments, ITeX, Siemens, Centillium, Silicon Automation Systems, and Nortel, proposed to use "a 180° phase shift" of the superframe synchronization symbol, *i.e.*, an inverted sync symbol, to coordinate a transmitter's and receiver's transition from an idle/quiescent mode to a higher-power (normal data) mode. *See, e.g.*, T1E1.4/98-279R1, CI-051, NF-059, WH-077, AB-080.

201. In 2000, Aware proposed that the transmitter would send "an inverted sync symbol as a flag" to tell the receiver when to change its configuration for SRA so that the transition would be synchronous with the transmitter's change. D.741, pp. 4-6. Aware proposed that the same inverted sync symbol could be used to coordinate transitions to and from a lower-power mode. *See, e.g.*, *id.* at pp. 6-8.

202. In 2000, Nortel proposed that the transmitter would "invert[] the current state of the sync symbol" each time a SRA was to be implemented to reduce the likelihood of the receiver not implementing the SRA synchronously with the transmitter. HC-071R1.

203. And, by no later than September 24, 2001, Syed Abbas and Guozhu Long had described, in a United States patent application, the use of a sync flag in "the 2nd generation ADSL standards" for OLR. *See* U.S. Patent Pub. No. 2002/0080867 ("Abbas"), ¶ [0007] ("In general, features such as bit swapping, rate adaptation, and bandwidth repartitioning techniques all require changes to a number of modulation parameters. Collectively these are also known as On Line Reconfiguration (OLR). The 2nd generation ADSL standards provide for an OLR protocol that allows a receiver to initiate any of the above mentioned changes through an OLR message sent over the modem overhead channel. If the proposed changes are not acceptable to the transmitter, the transmitter sends a NAK (negative acknowledge) message. Otherwise, the transmitter sends a sync flag that signals the proposed reconfiguration changes are acceptable and will take effect at a predetermined well defined time after the sync flag occurs."); *id.* at ¶ [0054] (describing 510 versions of shifted REVERB pseudo-random binary sequence and their inverted versions for use in signaling "e.g., entry and exit from Qmode or online modem reconfiguration").

204. Thus, several years before the '835 and '162 patents' priority date, the use of an inverted synchronization symbol as a flag transmitted by the transmitter to coordinate changes to transceivers' configurations during Showtime was known. This idea was no longer patentable as of the '835 and '162 patents' priority date.

b. Using a common protocol to coordinate multiple types of changes to transceiver configurations

205. Starting in 1998, 3Com, later joined by other companies such as Texas Instruments, ITeX, Siemens, Centillium, Silicon Automation Systems, and Nortel, proposed to augment the DRA protocol defined in T1.413 Issue 2 and G.992.1, which allowed FEC and interleaver parameters to be adjusted during Showtime, so that it could also be used to coordinate transitions between full-power mode and a lower-power mode during Showtime. *See, e.g.*, CI-051, NF-059, WH-077, T1E1.4/98-279R1.

206. In 2000, Catena Networks observed that "any rate adaptation or rate repartitioning" protocol could be used to coordinate framing changes, e.g., changes to "FEC and Mux Data frame parameters." BA-053, p. 3.

207. As described above, G.992.3 also defines a common OLR protocol to support bit swapping, DRR, and SRA. *See, e.g.*, G.992.3, § 10.2.

208. Thus, by at least 2000, the idea of using a common protocol to modify transceivers' configurations during Showtime was known in the art. Consequently, this idea was no longer patentable as of the '835 and '162 patents' priority date.

c. Modifying FEC and/or interleaver parameters during Showtime

209. As explained above, the DRA protocol standardized in both T1.413 Issue 2 and G.992.1 allowed FEC and interleaver parameters to be modified during Showtime. T1.413 Issue 2, Annex K; G.992.1, Appendix II.

210. Multiple contributions to SG15/Q4 also described and/or proposed to modify FEC and/or interleaver parameters during Showtime.

211. In June of 2000, Catena Networks described adaptive framing in the fast path to modify "the FEC and Mux Data Frame parameters" during Showtime to permit rate adaptation and repartitioning. BA-053, p. 1.

212. In 2000, Alcatel pointed out that SRA could be used to make FEC and interleaver configuration changes, and that these changes could be seamless under some conditions. HC-049, p. 1.

213. Thus, by at least 2000, the idea of making changes to the FEC and/or interleaver configuration during Showtime was known in the art. This idea was no longer patentable as of the '835 and '162 patents' priority date.

d. Making changes to transceiver parameters at FEC codeword boundaries

214. As explained above, G.992.1 and T1.413 Issue 2 both disclosed that if ADSL transceivers operate with the mandatory values of S (the number of DMT symbols per FEC codeword), the DRA protocol always resulted in changes being made at FEC codeword boundaries. G.992.1, § II.6. This aspect of the protocol allowed continued operation without the FEC mechanism needing to be reset. *Id.*

215. Several standards contributions also described or proposed making changes to transceiver parameters during Showtime at FEC codeword boundaries.

216. The Orckit proposal maintained the alignment of DMT symbols and FEC codewords from G.992.1 and proposed to use the DRA protocol of G.992.1. *See, e.g.*, FI-133, D.817. Accordingly, the Orckit proposal relied on making changes to FEC and/or interleaver parameters only on FEC codeword boundaries.

217. In 1999, Alcatel recognized that forcing transitions between a lower-power mode and a higher-power mode to occur only at the boundaries of FEC codewords would be desirable because it would allow seamless transitions and would simplify the implementation. *See, e.g.*, PO-038R1.

218. In 2000, Catena observed that if FEC codewords are not aligned with DMT symbols (as eventually became the case in G.992.3), configuration changes must be implemented on a FEC codeword boundary in order for those changes to be seamless. BA-053, p. 3.

219. Also in 2000, Alcatel stated again that changes to FEC and/or interleaver parameters "should coincide with RS codeword and/or Mux Data Frame boundary," and that "[o]nly a swap at that point can change the RS config[uration] param[eters]." HC-049, p 1.

220. Thus, by 2000, the idea of changing FEC and/or interleaver parameter values at FEC codeword boundaries was known in the art. Consequently, this idea was no longer patentable as of the '835 and '162 patents' priority date.

VIII. THE '835 AND '162 PATENTS

221. The '835 and '162 patents purport to have a priority date of March 3, 2004, which is well after all of the developments described above.

222. In the "Background" section, the '835 and '162 patents describe "a need in ADSL and VDSL systems to provide robust error-free performance in the presence of high, real-world impulse noise." '835 patent, 1:60-62. The patents note that "[i]t is standard practice for communications systems to use interleaving in combination with Forward Error Correction (FEC) to correct the errors caused by impulse noise." *Id.* at 1:34-37. According to the patents, the number of impulse-noise-corrupted DMT symbols that can be corrected by a particular

and/or the implementation timing information 427 (step 520)." *Id.* at ¶ [0061]. Then in step 530, "[t]he reconfiguration parameters 423, ack/comply timing information 425 and the implementation timing information 427 are then sent to the transmitter 410 via the OAM channel 420." *Id.* at ¶ [0062]. In step 545, "[a]fter receiving the reconfiguration parameters 423, ack/comply timing information 425 and/or the implementation timing information 427 from the receiver 440, the transmitter determines whether to implement the request for reconfiguration of transceive parameters." *Id.* at ¶ [0064]. In step 510, if the transmitter 410 determines that it will approve the reconfiguration request, "it will wait the indicated amount of time (step 547) and then acknowledge the request and indicate intent to comply (step 555). In the preferred embodiment, the transmitter acknowledges the request by sending an ack/comply signal 433 to the receiver 410, perhaps in the form of the synch_flag, at the time designated by the receiver 410, perhaps as designated by the value of SFlgSf." *Id.* at ¶ [0065]. In step 550, "the receiver 440 determines whether or not the superframe corresponding to flag SFlgSf includes the synch_flag." *Id.* at ¶ [0068]. If the receiver finds the synch_flag, "both the transmitter 410 and the receiver 440 will wait (in steps 565 and 560, respectively) before transmitting with the new reconfiguration transceive parameters." *Id.* at ¶ [0069].

B. The '835 Patent

1. Claim 10 Is Anticipated and/or Rendered Obvious by G.992.1

349. In my opinion, G.992.1 anticipates and/or renders obvious claim 10 of the '835 patent because G.992.1 discloses and/or renders obvious each limitation of claim 10.[8]

[8] T1.413 Issue 2, published in 1998, has disclosures substantively identical to those identified in this section for G.992.1. Accordingly, T1.413 Issue 2 also anticipates and/or renders obvious claims 8 and 10 for the same reasons that G.992.1 anticipates and/or renders obvious claims 8 and 10.

350. I provided an overview of G.992.1 above. *See supra*, § X.A.1. I incorporate that discussion by reference here.

a. "An apparatus configurable to adapt forward error correction and interleaver parameter (FIP) settings during steady-state communication or initialization"

351. G.992.1 discloses "[a]n apparatus configurable to adapt forward error correction and interleaver parameter (FIP) settings during steady-state communication or initialization."

352. Appendix II of G.992.1, entitled "Dynamic (on-line) Rate Adaptation," describes a mechanism called dynamic rate adaptation (DRA), which can be executed by the ATU-C and ATU-R and "allows reconfiguration of the modem during Showtime" without "a lengthy restart to reconfigure the modem." G.992.1, § II.1. G.992.1 explains that "[t]he purpose of this DRA mechanism is not to provide 'on-the-fly' Rate Adaptation, where the modem configuration would change continuously, tracking the slightest variation of the line conditions without affecting the user-traffic, but rather to allow for occasional changes, which would involve service interruption of the order of tens of milliseconds." *Id.*

353. The described DRA "is a mechanism that during ShowTime, without the need to restart: Allows rate modifications (up and downgrades) for both US and DS." *Id.* at § II.1.1. G.992.1 notes that "[r]ate modification implies more than just bit-rate but also FEC and Interleaving settings." *Id.*

354. The protocol "expands the AOC message set" and "define[s] new DRA-AOC messages" that allow the ATU-C to "propose to the ATU-R a new rate configuration." *Id.* at § II.2.1. Other messages are used by the transceivers to "exchange configuration information – if the proposal is accepted by ATU-R." *Id.* Finally, the DRA-AOC messages allow the transceivers to "initiate and synchronize a swap to the new rate configuration." *Id.*

362. G.992.1 discloses that the described transceivers include a processor. *See, e.g.*, § 5.1.1 ("the first incoming bit (outside world MSB) shall be the first processed bit inside the ADSL (ADSL LSB)."); *id.* at § 5.1.2 ("the first incoming bit (outside world MSB) will be the first processed bit inside the ADSL (ADSL LSB), and the CLP bit of the ATM cell header will be carried in the MSB of the ADSL frame byte (i.e. processed last)."); *id.* at § 5.2.1 ("the first incoming bit (outside world MSB) will be the first processed bit inside the ADSL (ADSL LSB)."); *id.* at § 5.2.2 ("the first incoming bit (outside world MSB) will be the first processed bit inside the ADSL (ADSL LSB), and the CLP bit of the ATM cell header will be carried in the MSB of the ADSL frame byte (i.e. processed last)."); *id.* at Figure 5-1 (showing that ATU-C includes "DAC and Analogue processing"); *id.* at Figure 5-2 (same); *id.* at Figure 5-3 (showing that ATU-R includes "DAC and Analogue processing"); *id.* at Figure 5-4 (same); *id.* at § 7.4 ("the first incoming bit (outside world MSB) will be the first processed bit inside the ADSL (ADSL LSB)."); *id.* at § 8.4 (same); *id.* at § 9.2.3.1 ("This message tells the ATU-R to maintain the ATU-R EOC processor and any active ADSL EOC-controlled operations (such as latching commands) in their present state. . . . This message releases all outstanding EOC-controlled operations (latched conditions) at the ATU-R and returns the ADSL EOC processor to its initial state.").

363. To the extent it is determined that G.992.1 does not explicitly or inherently disclose a processor, the inclusion of a processor in an ADSL transceiver would have been obvious to a person having ordinary skill in the art as of the priority date of the '835 patent. *See, e.g.*, Charles K. Summers, *ADSL Standards, Implementation, and Architecture*, CRC Press (1999), p. 74 ("The interface chip is the chip at the heart of telecommunications equipment. The interface chip deals with the physical medium in the manner necessary to perform the

physical layer protocol. In the case of ADSL, this means performing the functions of three major functional blocks: Digital Interface (DI), Digital Signal Processing (DSP), and Analog Interfaces (AI)."); *id.* at p. 84 ("[A] microprocessor needs to access information within an interface chip. The memory map indicates that all I/O performed between 0x4000 to 0x43FF (a range of 1024 identifiers) will be routed between the requesting device and the interface chip. The 'address decoding logic' indicates that when this address is written onto the address bus, a sequence of instructions will commence. . . . The net effect will be that a write of data to address 0x4204 will go to the interface chip and that, because of the specific location within the address range, some type of behavior is expected as a result of the write."); *id.* at p. 87 ("A Low-Level Driver (LLD) acts as the software interface between a semiconductor device and the controlling software. . . . The LLD must be able to interpret primitives from, and send primitives to, the software module (or modules) with which it is associated."); *id.* at p. 94 ("Our discussion on LLDs indicated that that ADSL chipset interface was needed to create an interface between the software (running on a controlling general-purpose microprocessor) and the hardware.").

364. Moreover, and as discussed elsewhere herein, in my experience, every DSL transceiver includes some kind of processor (e.g., a digital signal processor, microprocessor, etc.).

365. Accordingly, G.992.1 discloses "a transceiver, including a processor, configurable to" perform functions.

c. "transmit a signal using a first FIP setting"

366. G.992.1 discloses that each of the ATU-C and ATU-R is configurable to transmit a signal using a first FIP setting. *See, e.g.*, G.992.1, §§ 7.6, 8.6, Table 7-7, Table 8-3.

367. G.992.1 sets forth minimum required downstream and upstream FEC and interleaving capabilities for both the ATU-C and ATU-R, including all of the values of R and

slow and it does not guarantee error free PCM data. . . . The G.992.1 Appendix II definition (Section II.6) requires that the SFR value sent by the ATU-C must be at least 47 superframes (800 msec) greater than the present superframe counter value (t=0). This delay is to guarantee sufficient time for reception and reconfiguration calculations to be ready for the DRA transition. However, for DRR events, the reconfiguration is significantly lower in complexity since no bits and gains tables are modified. Furthermore, the delay associated with 47 superframes makes any dynamic bandwidth allocation for PCM voice violate telephony call setup/tear down timing requirements."). Therefore, a person having ordinary skill in the art would have considered whether the DRA protocol of G.992.1 could be made faster for applications such as DRR for telephone call setup and tear-down.

391. As I explained previously, during the work on idle mode (later to be known as Qmode or quiescent mode), the use of an inverted synchronization symbol with the G.992.1 DRA protocol was proposed to allow fast exits from idle mode. *See, e.g.*, CI-051, p. 3 ("In an effort to minimize the required changes, 3Com suggests that two new commands be added to the Dynamic Rate Adaptation command set: DRA_Idle_Request: Generated at either the ATU-C or ATU-R, this command would signal a request to enter idle mode at a particular superframe boundary, as specified by the superframe counter. DRA_Idle_Reply: Generated at the receiver, this command would be a response to the DRA_Idle_Request. . . . For exiting idle mode, 3Com proposes that the final idle mode symbol be shifted 180° with respect to the other idle mode symbols (with the exception of the pilot carrier which may be required for timing recovery). This can be initiated by either the ATU-C or ATU-R. Upon detection, the receiver assumes the next symbol is the first symbol of the first superframe following the idle period interruption. The latency of this proposal is one DMT symbol and thus has no effect on the data latency and

imposes no additional buffering requirements."); *id.* at p. 2 ("3Com proposes that the superframe synchronization symbol be used as the idle symbol.").

392. Accordingly, the use of an inverted sync symbol, which the '835 patent states is a flag signal, ('835 patent, 12:29-31), had already been proposed in 1998 as a modification to improve the speed at which changes to transceiver configurations could be made. Thus, modifying the DRA protocol of G.992.1 to use an inverted sync symbol instead of the DRA_Swap_Reply message to improve the speed of dynamic rate reconfigurations for voice applications would have been obvious. G.992.1 disclosed that the superframe boundaries coincide with the FEC codeword boundaries at least every $4 \times N - 1$ superframes, where N is an integer number. Thus, it would have been obvious for the ATU-R to transmit an inverted sync symbol instead of the DRA_Swap_Reply message to acknowledge the DRA_Swap_Request message from the ATU-C. This inverted sync symbol would confirm to the ATU-C that the switch to a new FIP setting will occur at the superframe boundary specified by the SFR value sent by the ATU-C in the DRA_Swap_Request message.

393. Thus, G.992.1 discloses and/or renders obvious a transceiver that is configurable to transmit a flag signal.

e. <u>"switch to using for transmission, a second FIP setting following transmission of the flag signal"</u>

394. G.992.1 discloses that both the ATU-C and ATU-R switch to using for transmission a second FIP setting following transmission of the flag signal.

395. G.992.1 discloses that the DRA procedure defined in Appendix II "allows reconfiguration of the modem during Showtime" without "a lengthy restart to reconfigure the modem." G.992.1, § II.1. The described DRA "is a mechanism that during ShowTime, without the need to restart: Allows rate modifications (up and downgrades) for both US and DS." *Id.* at

(BS, SRA, DRR)."); *id.* ("The OLR protocol is initiated by an ATU requesting to reconfigure the ATU's receiver parameters only. This means that the downstream parameters can only be reconfigured by an OLR initiated by an ATU-R and that the upstream parameters can only be reconfigured by an OLR initiated by ATU-C."); *id.* at pp. 5-10 (describing OLR protocol, indicating both ATU1 and ATU2 can both transmit and receive). A person having ordinary skill in the art would have recognized that the transmitting and receiving portions of an ATU-C (or ATU-R) share common circuitry (e.g., an interface to the twisted pair).

429. As a person having ordinary skill in the art would also have recognized, the ATU-C and ATU-R described in SC-060 include a processor. *See, e.g.*, *id.* at p. 1 ("As many Bit Swaps may be required in a short period of time, a key feature of the OLR protocol should be to process Bit Swap OLRs in a minimum amount of time."); p. 6 ("Instead of having to search for a potential Synch Flag during every sync symbol, the requesting ATU can utilize the sync symbol periods of time up to the scheduled sync symbol to perform other critical processing tasks."). As of the priority date of the '835 patent, it was common for ATU-Cs and ATU-Rs to include a processor.

430. Accordingly, SC-060 discloses a transceiver, including a processor, that is configurable to perform various functions.

c. "transmit a signal using a first FIP setting"

431. SC-060 discloses a transceiver configurable to transmit a signal using a first FIP setting.

432. SC-060 discloses both an ATU-C and an ATU-R. *See, e.g.*, SC-060, § 3. As would have been understood by a person having ordinary skill in the art as of the '835 patent's priority date, ATU-Cs and ATU-Rs have always been capable of transmitting signals using both FEC and interleaving. *See, e.g.*, T1.413 Issue 1, §§ 6.4.1, 6.4.2; *see also supra*, §§ VII.D,

X.B.1.c. Indeed, the '835 patent itself admits in its background section that "[i]t is standard practice for communications systems to use interleaving in combination with Forward Error Correction (FEC) to correct errors caused by impulse noise," and that it was known in ADSL that FEC and interleaving parameters include N, K, R, and D, "where N is the codeword size in bytes, R is the number of parity (or redundancy) bytes in a codeword, D is the interleaver depth in number of codewords," and "K is the number of information bytes in a codeword." '835 patent, 1:34-37, 2:12-16. Therefore, SC-060 discloses to skilled artisans a transceiver configurable to transmit a signal using a first FIP setting.

433. Furthermore, SC-060 discloses that the FEC codeword size could be modified using the disclosed OLR protocol. A person having ordinary skill in the art would have understood from this disclosure that the ATU-C and ATU-R of SC-060 would include at least FEC block coding, and that the FEC codeword size would be the sum of the number of payload bytes and the number of redundancy bytes because it was known that N = K + R for the Reed-Solomon coding used in ADSL.

434. Therefore, a person having ordinary skill in the art would have understood that the "ATU-C" and "ATU-R" described in SC-060 would be configurable to transmit a signal using a first FIP setting.

435. Accordingly, SC-060 discloses transceivers that are configurable to transmit a signal using a first FIP setting.

d. <u>"transmit a flag signal"</u>

436. As construed by the Court, a flag signal is a "signal used to indicate when an updated FIP setting is to be used (the signal does not include the FEC codeword counter value upon which the updated FIP setting is to be used)." SC-060 discloses a transceiver that is configurable to transmit a flag signal, as construed by the Court.

indicate when an updated FIP setting is to be used" that "does not include the FEC codeword counter value upon which the updated FIP setting is to be used."

441. In addition, I understand that TQ Delta has identified the syncflag of G.993.2, which is used to signal on-line reconfiguration transitions, as allegedly meeting this element of claim 8 (and, therefore, claim 10). The Synch Flag described by SC-060 has the same characteristics as the syncflag in G.993.2. Accordingly, to the extent it is determined that the syncflag of G.993.2 meets this element, SC-060 also discloses this element.

442. Accordingly, SC-060 discloses a transceiver configurable to transmit a flag signal.

e. "switch to using for transmission, a second FIP setting following transmission of the flag signal"

443. SC-060 discloses that both the ATU-C and ATU-R switch to using for transmission a second FIP setting following transmission of the flag signal.

444. SC-060 discloses a transceiver that is configurable to switch to using for transmission a second transceiver setting. *See, e.g.*, SC-060, 1-2 (disclosing use of OLR to modify bits and fine gain parameters, number of octets per bearer mux data frame, data rate of connection, etc.). SC-060 also states that it is possible to modify "the FEC Codeword size" and other PMS-TC parameters in DRR and SRA transactions. *Id.* at p. 2. As a person having ordinary skill in the art would have understood, PMS-TC parameters include FEC and interleaving parameters. *See, e.g.*, G.993.1, § 8 ("The PMS-TC sublayer provides transmission medium specific TC functions, such as framing, forward error correction (FEC), and interleaving."). Moreover, the claimed FIP settings include the FEC Codeword size, denoted as N in the '835 patent.

445. SC-060 also discloses that the switch to a second transceiver setting follows transmission of the flag signal. For example, SC-060 discloses that after transmitting the Synch Flag, the transmitting transceiver "shall start using the new configuration effective the first PMD symbol of the PMD superframe count equal to (SFlgSf + 1 + SfDly) modulo 256," where "SfDly is a non-negative integer less than 4 and is approximately the number of PMD superframes that the receiver will have advance knowledge of an impending OLR reconfiguration before the new configuration must become effective." *Id.* at p. 4. Figure 2 of SC-060, copied below, illustrates a successful completion of the OLR protocol and indicates that the reconfiguration (*i.e.*, the switch to the second FIP setting) occurs after transmission of the flag signal ("SYNCH FLAG").

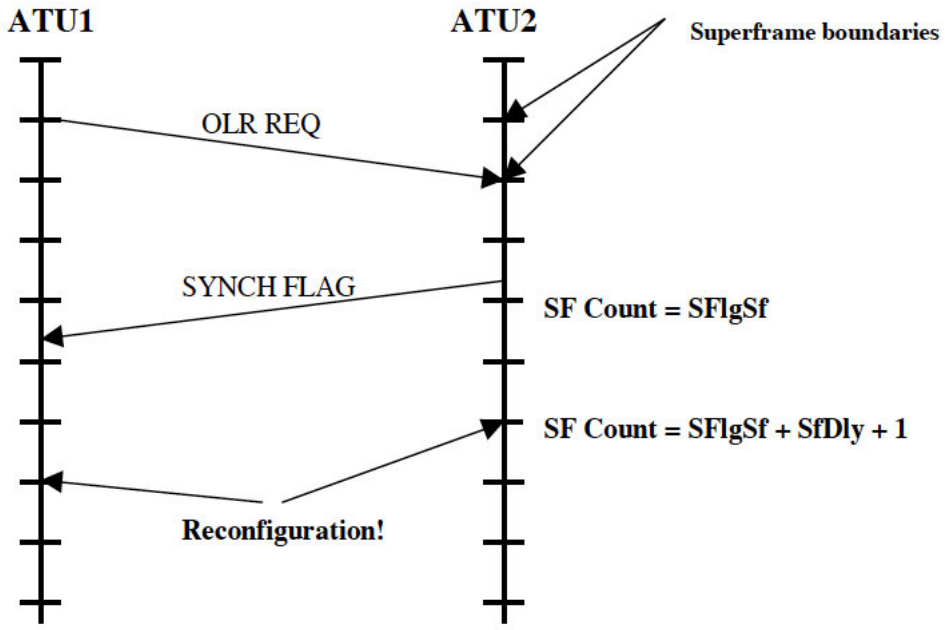

Figure 2: a successful completion of the OLR protocol

See also id. at p. 5 (explaining that on the first DMT symbol of the Superframe with count equal to (SFlgSf + 1 + SfDly) modulo 256, both ATUs use new parameters).

446. As explained above with respect to the preamble and the flag signal, to the extent that SC-060 does not disclose using the OLR protocol to reconfigure FIP settings, it would have been obvious to a person having ordinary skill in the art to do so.

447. Thus, it would have been obvious to a person having ordinary skill in the art to extend the SC-060 OLR protocol to cover modifications to other PMS-TC parameters (e.g., K, R, N, D) such that the ATU2 illustrated in Figure 2 would switch to using for transmission, a second FIP setting following transmission of the Synch Flag, i.e., the flag signal, and that the ATU1 would also switch to using for reception, a second FIP setting following reception of the Synch Flag, i.e., the flag signal.

448. Thus, SC-060 discloses that the ATU-C and ATU-R switch to using for transmission, a second FIP setting following transmission of the flag signal.

f. "wherein: the first FIP setting comprises at least one first FIP value"

449. SC-060 discloses that the first FIP setting comprises at least one first FIP value.

450. SC-060 discloses that seamless rate adaptation may be used to modify "the FEC Codeword size" and other PMS-TC parameters. SC-060, p. 2. As would have been understood by a person having ordinary skill in the art as of the '835 patent's priority date, the FEC codeword size is a FIP parameter, specifically a FEC parameter, having a numerical value. Furthermore, it was known as of the '835 patent's priority date that PMS-TC parameters include FEC and interleaver parameters. *See, e.g.*, G.993.1, § 8 ("The PMS-TC sublayer provides transmission medium specific TC functions, such as framing, forward error correction (FEC), and interleaving.").

451. Furthermore, a person having ordinary skill in the art would have known that ATU-Cs and ATU-Rs have always supported both FEC and interleaving. *See, e.g.,* T1.413 Issue

1, §§ 6.4.1, 6.4.2; G.992.1, §§ 7.6, 8.6, Table 7-7, Table 8-3; *see also supra*, §§ VII.D, X.B.1.f. Thus, a person having ordinary skill in the art would have understood the ATU-C and ATU-R of SC-060 to transmit and receive signals using FIP settings, which would include at least one first FIP value.

452. Accordingly, SC-060 discloses that the first FIP setting comprises at least one first FIP value.

g. "the second FIP setting comprises at least one second FIP value, different than the first FIP value"

453. SC-060 discloses that the second FIP setting comprises at least one second FIP value different than the first FIP value.

454. SC-060 discloses that the ATUs are configurable to perform various types of on-line reconfiguration, including modifying "the FEC codeword size" and, generally, "PMD, PMS-TC, and TPS-TC parameters." SC-060, p. 2. Because, as was known as of the '835 patent's priority date, PMS-TC parameters include forward error correction and interleaving parameters (*see, e.g.*, G.993.1, § 8), and the ATU-C and ATU-R of SC-060 are configurable to make changes to PMS-TC parameters, SC-060 discloses that the second FIP setting comprises at least one second FIP value, different than the first FIP value.

455. SC-060 also discloses that a seamless rate adaptation can be used to modify "the FEC Codeword size." SC-060, p. 2. As would have been understood by a person having ordinary skill in the art as of the '835 patent's priority date, following such modification, the FEC codeword size, which is a FEC parameter having a value, would have a different numerical value than the one it had before the modification.

456. Thus, SC-060 discloses that the second FIP setting comprises at least one second FIP value, different than the first FIP value.

effective." SC-060, § 3.4. Figure 1 of SC-060, copied below, illustrates the timing of the switch to the new configuration.

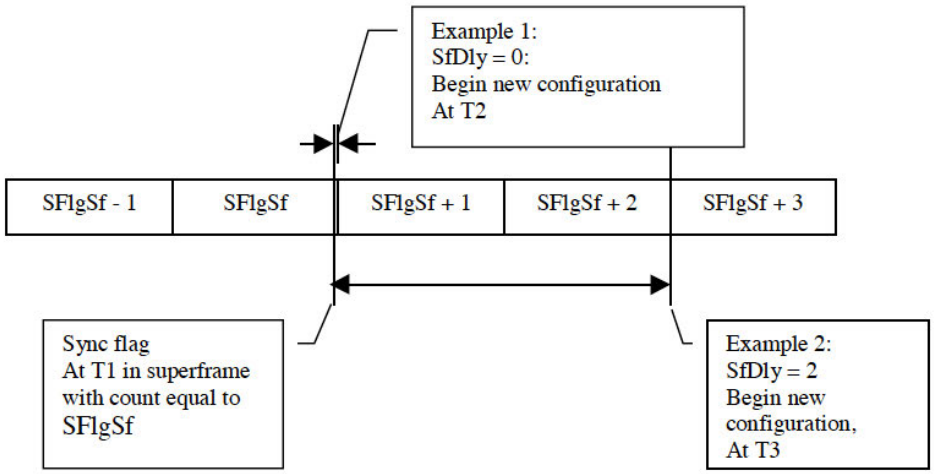

Notes: Each block is a PMD superframe and the text within block is the superframe count. Two examples shown. Example 1 with SfDly = 0 giving receiver less than T2 – T1 or less than 250 µSec advance notice. Example 2 with SfDly = 2 giving receiver T3 – T1 or about 2 PMD superframes.

Figure 1: Timing of OLR Exchange.

461. A person having ordinary skill in the art would have been intimately familiar with G.992.1 as of the '835 patent's priority date and would have understood that it would be necessary for the ATU transmitting the OLR Request message to choose values of SFlgSf and SfDly such that the beginning of the first PMD symbol of the PMD superframe count equal to (SFlgSf + 1 + SfDly) modulo 256 coincides with an FEC codeword boundary whenever a FIP setting is being modified. As I have explained, it was known that making changes on FEC codeword boundaries was necessary to avoid a reset of the FEC mechanism. *See, e.g.*, G.992.1 at § II.6; HA-049; D.817; PO-038R1; BA-053. The positions of the FEC codeword boundaries are always known to the transmitter (which is generating the FEC codewords) and the receiver (which is receiving and decoding the FEC codewords).

462. A person having ordinary skill in the art would also have understood that even if the FEC codeword boundaries are not aligned with DMT symbol boundaries, the FEC codeword boundaries and DMT symbol boundaries still coincide periodically. *See, e.g.*, BA-

053, p. 3 ("worst case periodicity for achieving DMT symbol – codeword alignment if they are not byte or codeword aligned is 2040 DMT symbols or 0.51 seconds."); D.834, p. 3 ("The boundaries of the Constellation Encoder Input Data Frame (i.e., symbol boundary) and the RS codeword are aligned at least at the beginning of SHOWTIME and further every (8*NFEC) symbols. The alignment occurs therefore at least every 8*255=2040 data symbols or 0.51 seconds."). As a person having ordinary skill in the art would have recognized, the FEC codeword boundaries aligning periodically with the FEC codeword boundaries means that the switching to an updated FIP setting is effective on the boundary of a FEC codeword where the position of the boundary of each codeword is known prior to the switching.

463. Furthermore, "the switching occurs on a pre-defined forward error correction codeword boundary following the flag signal" would have been obvious in view of the disclosures of SC-060, at least because it was known that changing FEC and/or interleaver parameters on a FEC codeword boundary would avoid an explicit reset of the FEC encoder and decoder. *See* G.992.1, § II.6; *see also* G.992.3, § 7.11.1.1 (listing "the Interleaved FEC data frame boundary, FEC data frame boundary, Mux Data Frame boundary, and the PMD symbol boundary" as locations for on-line reconfiguration).

464. In addition, a person having ordinary skill in the art would have understood that making modifications to FIP settings at FEC codeword boundaries would be an obvious choice because of the way FEC encoders create codewords. The transmitter generates the FEC codewords, and the FEC codewords are the smallest "units" the FEC encoder can produce because the redundancy bytes are calculated from all of the payload bytes in the codeword. Therefore, the transmitter always knows the locations of the FEC codeword boundaries and is configurable to switch from a first FIP setting to a second FIP setting on any of these known

boundaries. Indeed, on an FEC codeword boundary is the only place where the transmitter can change FEC parameters. *See, e.g.*, HA-049, p. 1 ("Should SRA messages also include Mux Data Frame, RS Codeword and interleaver configuration changes? . . . The swap superframe boundary should then coincide with RS codeword and/or Mux Data Frame boundary. For RS, that is worst case every 4.3 sec (if L even). Only a swap at that point can change the RS config param.").

465. Thus, it would have been obvious to a person having ordinary skill in the art in extending the SC-060 protocol to cover changes to FIP settings to restrict the values of the scheduling parameters, SFlgSf and SfDly, such that the switch to the new FIP setting would always occur on a pre-defined FEC codeword boundary.

466. Accordingly, SC-060 discloses or renders obvious that the switching occurs on a pre-defined forward error correction codeword boundary following the flag signal.

i. "wherein a first interleaver parameter value of the first FIP setting is different than a second interleaver parameter of the second FIP setting"

467. SC-060 discloses and/or renders obvious that a first interleaver parameter value of the first FIP setting is different than a second interleaver parameter of the second FIP setting.

468. SC-060 discloses that both DRR and SRA can be used to modify PMS-TC parameters, which, as a person having ordinary skill in the art would have understood, include interleaver parameters, such as the interleaver depth. *See, e.g.*, G.992.3, § 7.2 ("In addition to transport functionality, the ATU transmit PMS-TC function also provides procedures for: • scrambler; • insertion of redundancy for Reed-Solomon-based forward error correction; • insertion of checksums for block based error detection; and • interleaving of data frames to spread the effect of impulsive impairments on the U interface."); *id.* at § 7.5 (listing PMS-TC control parameters, which include "D_p," defined as "[t]he interleaving depth in the latency path

function #p."); *see also, e.g.*, G.993.1, § 8 ("The PMS-TC sublayer provides transmission medium specific TC functions, such as framing, forward error correction (FEC), and interleaving.").

469. Accordingly, SC-060 discloses and/or renders obvious that a first interleaver parameter value of the first FIP setting is different than a second interleaver parameter of the second FIP setting.

470. Consequently, claim 10 is anticipated and/or rendered obvious by SC-060.

3. Claim 10 Is Rendered Obvious by G.992.1 in Combination With SC-060

471. In my opinion, G.992.1 in combination with SC-060 renders claim 10 of the '835 patent obvious because G.992.1 and SC-060 together disclose each limitation of claim 10. In addition, a person having ordinary skill in the art as of the '835 patent's priority date would have been motivated to combine the disclosures of G.992.1 with those of SC-060 as recited in claim 10.

472. I provided an overview of G.992.1 above. *See supra*, § X.A.1. I incorporate that discussion by reference here.

473. I provided an overview of SC-060 above. *See supra*,§ X.A.2. I incorporate that discussion by reference here.

a. "An apparatus configurable to adapt forward error correction and interleaver parameter (FIP) settings during steady-state communication or initialization"

474. G.992.1 in combination with SC-060 discloses "[a]n apparatus configurable to adapt forward error correction and interleaver parameter (FIP) settings during steady-state communication or initialization."

buffer is used. In this case, the aoc messages used to coordinate reconfigurations during Showtime, including DRA, are transmitted in the fast path and are not protected by interleaving. As a result, when only the fast buffer is used, the reconfiguration protocol of G.992.1 is less robust (more prone to errors) than it would otherwise be.

531. Thus, a person having ordinary skill in the art would have been motivated to improve the robustness of the DRA protocol, and specifically the acknowledgment of the timing of the new configuration, to provide improved "protection against transmission errors." *See, e.g.*, ITU-T SG15/Q4 Contribution BI-063, G.gen: A Fast Approach for DRR/DRA/Bitswap Event Signaling, ("BI-063"), § 1 (describing need for "features that help the robustness and performance of the modems" and stating that "there is a need for a sufficiently fast, reliable and robust signaling method to allow dynamic rate repartitioning," which "can be used for bit swapping and DRA parameter changes as well and would increase the utility of these methods"). Specifically, a person having ordinary skill in the art would have sought ways to improve the robustness of the DRA protocol regardless of whether transceivers operate with dual latency or in the reduced overhead mode. To identify ways to improve the DRA protocol of G.992.1, a person having ordinary skill in the art would have looked to the contributions to the relevant ITU-T working group, namely the SG15/Q4 working group, and these contributions would have included SC-060.

532. SC-060 proposes a unified protocol for multiple types of on-line reconfiguration (OLR). SC-060, § 2.1. SC-060 explicitly discloses modifying PMS-TC parameters in OLR transactions. *Id*. As a person having ordinary skill in the art as of the '835 patent's priority date would have understood, PMS-TC parameters include FEC and interleaver parameters, such as the FEC codeword size, the number of information bytes per FEC codeword, the number of

redundancy bytes per FEC codeword, and the interleaver depth. *See, e.g.*, G.992.3, § 7.2 ("the ATU transmit PMS-TC function also provides procedures for: . . . interleaving of data frames to spread the effect of impulsive impairments"); *id.* at § 7 (entitled "Physical Media Specific Transmission Convergence (PMS-TC) function" and including subsection 7.7.1.5, entitled "Interleaver" and indicating the interleaver is characterized by, among other things, an interleaver depth D); *id.* at § 7.5 (indicating that number of payload bytes per FEC codeword, number of redundancy bytes per FEC codeword, and interleaving depth are PMS-TC parameters). SC-060 teaches that in addition to the bit swapping, dynamic rate repartitioning, and seamless rate adaptation OLR types already defined, "the FEC Codeword size could also be modified." SC-060, p. 2. Accordingly, SC-060 discloses a generic OLR protocol that "obviously" can be used for a "wide variety of potential OLR applications." *Id.*at § 2.1, Abstract.

533. In the SC-060 protocol, the receiving ATU ("ATU1") sends an OLR Request message to the transmitting ATU ("ATU2"). *Id.* at §§ 3.1, 3.5. The OLR Request message includes proposed modified transceiver settings, as well as two additional parameters, Synch Flag Superframe (SFlgSf) and Superframe Delay (SfDly). *Id.* The value of SFlgSf instructs ATU2 when to send an OLR Acknowledgement to ATU1. *Id.* at §§ 3.1, 3.4. The value of SfDly tells ATU2 how long after sending the OLR Acknowledgement it should wait before switching to the reconfigured transceiver settings. *Id.* at §§ 3.1, 3.4, 3.5.

534. SC-060 discloses that the OLR Acknowledgement is conveyed by ATU2 sending a Synch Flag in place of a synch symbol as the 69th frame of the superframe corresponding to SFlgSf. *Id.* at §§ 3.3, 3.5. SC-060 describes the Synch Flag as the inverse of the synch symbol. *Id.* at § 3.3. ATU2 sends the OLR Acknowledgement (the Synch Flag) after

the PMD layer, in the superframe corresponding to SFlgSf, as taught by Wunsch, and (c) the transceivers both switching to the new configuration after (SFlgSf + 1 + Dly) modulo 256 superframes, as taught by Wunsch. In this manner, the Synch Flag, sent by the ATU-R, would be used to indicate to the ATU-C when to use the updated FIP setting (the new upstream settings for reception), and to indicate when the ATU-R will use the new upstream settings for transmission.

562. Accordingly, G.992.1 in combination with Wunsch discloses "[a]n apparatus configurable to adapt forward error correction and interleaver parameter (FIP) settings during steady-state communication or initialization."

b. "a transceiver, including a processor, configurable to:"

563. In my opinion, G.992.1 in combination with Wunsch discloses a transceiver, including a processor. As explained in sections below, the transceiver is configurable to perform various functions, including those functions required by claim 10.

564. As I explained above (*see supra*, § X.B.1.b), G.992.1 discloses a transceiver including a processor that is configurable. I incorporate that discussion by reference here.

565. Wunsch also discloses a transceiver including a processor that is configurable. *See, e.g.*, Wunsch, ¶ [0005] ("Referring to FIG. 6, the receiver in transceiver B monitors some measure of the relative error performance of each of the received carriers in a multi-carrier modulation such as is used in ADSL. If the receiver in transceiver B determines that a change is desirable, e.g. to decrease the overall error level, then transceiver B sends a request message that contains the proposed change to transceiver parameters to transceiver A, which sends an acknowledge/will comply (hereafter called ack/comply message) back to the transceiver B if it wishes to make the requested change."); *id.* at ¶ [0003] ("Many protocols have been suggested to help more efficiently and accurately transmit and process data in the DSL systems.").

566. To the extent it is determined that Wunsch does not explicitly or inherently disclose a processor, the inclusion of a processor in the transceivers described in Wunsch would have been obvious to a person having ordinary skill in the art as of the priority date of the '835 patent. *See, e.g.*, Charles K. Summers, *ADSL Standards, Implementation, and Architecture*, CRC Press (1999), p. 74 ("The interface chip is the chip at the heart of telecommunications equipment. The interface chip deals with the physical medium in the manner necessary to perform the physical layer protocol. In the case of ADSL, this means performing the functions of three major functional blocks: Digital Interface (DI), Digital Signal Processing (DSP), and Analog Interfaces (AI)."); *id.* at p. 84 ("[A] microprocessor needs to access information within an interface chip. The memory map indicates that all I/O performed between 0x4000 to 0x43FF (a range of 1024 identifiers) will be routed between the requesting device and the interface chip. The 'address decoding logic' indicates that when this address is written onto the address bus, a sequence of instructions will commence. . . . The net effect will be that a write of data to address 0x4204 will go to the interface chip and that, because of the specific location within the address range, some type of behavior is expected as a result of the write."); *id.* at p. 87 ("A Low-Level Driver (LLD) acts as the software interface between a semiconductor device and the controlling software. . . . The LLD must be able to interpret primitives from, and send primitives to, the software module (or modules) with which it is associated."); *id.* at p. 94 ("Our discussion on LLDs indicated that that ADSL chipset interface was needed to create an interface between the software (running on a controlling general-purpose microprocessor) and the hardware.").

567. Moreover, and as discussed elsewhere herein, in my experience, every DSL transceiver includes some kind of processor (e.g., a digital signal processor, microprocessor, etc.).

568. Accordingly, G.992.1 in combination with Wunsch discloses "a transceiver, including a processor, configurable to" perform functions.

c. "transmit a signal using a first FIP setting"

569. G.992.1 in combination with Wunsch discloses a transceiver that transmits a signal using a first FIP setting and a transceiver that receives a signal using a first FIP setting.

570. As I explained above (*see supra*, § X.B.1.c), G.992.1 discloses transceivers configurable to transmit signals using a first FIP setting. I incorporate that discussion by reference here.

571. Wunsch also discloses a transceiver including a processor that is configurable to transmit signals using a first FIP setting. Wunsch discloses that "[t]he preferred embodiment of the invention improves the ability for ADSL equipment to reconfigure transceive parameters during normal 'showtime' operation." *See, e.g.*, Wunsch ¶ [0018]. As would have been understood by a person having ordinary skill in the art as of the '835 patent's priority date, ADSL equipment has always been capable of transmitting signals using both FEC and interleaving. *See, e.g.*, T1.413 Issue 1, §§ 6.4.1, 6.4.2; *see also supra*, § X.B.1.c. Indeed, the '835 patent itself admits in its background section that "[i]t is standard practice for communications systems to use interleaving in combination with Forward Error Correction (FEC) to correct errors caused by impulse noise," and that it was known in ADSL that FEC and interleaving parameters include N, K, R, and D, "where N is the codeword size in bytes, R is the number of parity (or redundancy) bytes in a codeword, D is the interleaver depth in number of codewords," and "K is the number of information bytes in a codeword." '835 patent, 1:34-37, 2:12-16. Therefore, Wunsch inherently discloses a transceiver configurable to transmit a signal using a first FIP setting.

572. Accordingly, G.992.1 in combination with Wunsch discloses a transceiver that transmits a signal using a first FIP setting and a transceiver that receives a signal using a first FIP setting.

d. "transmit a flag signal"

573. The term "flag signal" has been construed as a "signal used to indicate when an updated FIP setting is to be used (the signal does not include the FEC codeword counter value upon which the updated FIP setting is to be used)." Claim Construction Order, 91.

574. G.992.1 in combination with Wunsch discloses that the ATU-R transmits, and the ATU-C receives, a signal to indicate when an updated FIP setting is to be used, and this signal does not include the FEC codeword counter value upon which the updated FIP setting is to be used.

575. G.992.1 discloses that in order to propose a new configuration, the ATU-C sends the DRA_Configuration_Request message, which includes proposed values for the number of parity bytes per symbol in the downstream fast buffer (R_{fd}), the number of parity bytes per symbol in the downstream interleave buffer (R_{id}), the number of parity bytes per symbol in the upstream fast buffer (R_{fu}), the number of parity bytes per symbol in the upstream interleave buffer (R_{iu}), the number of downstream symbols per FEC codeword (S_d), the number of upstream symbols per FEC codeword (S_u), the downstream interleave depth in codewords (I_d), and the upstream interleave depth in codewords (I_u). G.992.1, § II.4.1. G.992.1 discloses that through a series of additional messages the transceivers agree to implement the reconfiguration. *Id.* at § II.4.2.

576. The ATU-C then sends the DRA_Swap_Request message "to inform the ATU-R about when to swap the rate." *Id.* at § II.6. A superframe reference number, SFR, identifies "around which superframe boundary the rate swap will occur." *Id.* The ATU-R acknowledges

f. "wherein: the first FIP setting comprises at least one first FIP value"

587. G.992.1 in combination with Wunsch discloses that the first FIP setting comprises at least one FIP value.

588. As I explained above (*see supra*, § X.B.1.f), G.992.1 discloses that the first FIP setting comprises at least one first FIP value. I incorporate that discussion by reference here.

589. Wunsch also discloses that the first FIP setting comprises at least one FIP value. Wunsch discloses that "[t]he preferred embodiment of the invention improves the ability for ADSL equipment to reconfigure transceive parameters during normal 'showtime' operation." *See, e.g.*, Wunsch, ¶ [0018]. As would have been understood by a person having ordinary skill in the art as of the '835 patent's priority date, ADSL equipment has always been capable of transmitting signals using both FEC and interleaving during "normal 'showtime' operation." *See, e.g.*, T1.413 Issue 1, §§ 6.4.1, 6.4.2; *see also supra*, § X.B.1.f. Therefore, a person having ordinary skill in the art would have understood that the ADSL transceivers described in Wunsch are configurable to transmit a signal using a first FIP setting, wherein the first FIP setting comprises at least one FIP value.

590. Accordingly, G.992.1 in combination with Wunsch discloses that the first FIP setting comprises at least one FIP value.

g. "the second FIP setting comprises at least one second FIP value, different than the first FIP value"

591. G.992.1 in combination with Wunsch discloses that the second FIP setting comprises at least one second FIP value, different than the first FIP value.

592. As I explained above (*see supra*, § X.B.1.g), G.992.1 discloses that the second FIP setting comprises at least one second FIP value, different than the first FIP value. I incorporate that discussion by reference here.

Showtime, which would result in a first interleaver parameter value of the first FIP setting being different from a second interleaver parameter value of the second FIP setting. As I have explained, making changes to interleaver parameter values during Showtime had been standardized in G.992.1. It was well known that changes to the interleaver depth could be made via on-line reconfiguration protocols. Accordingly, a person having ordinary skill in the art as of the '835 patent's priority date would have recognized that the disclosed protocols of Wunsch are not limited to transceivers performing bit swapping, but also extend to transceivers performing other types of on-line reconfigurations, such as those necessitated by changes in network traffic. *See, e.g.*, Wunsch, ¶ [0004] (referring to bit swapping as "one such protocol" used to make changes during operation); *id.* at ¶ [0005] (explaining that "other modifications to transceive parameters" in addition to bit swapping can be made). Such changes would include changes to the interleaver depth.

605. Furthermore, Wunsch contrasts its protocol with "existing procedures" and "potential reconfiguration procedures being discussed in [the] ITU-T standards group." *Id.* at ¶ [0018]. A person having ordinary skill in the art would have understood that these "existing procedures" included the DRA protocol of G.992.1, which allowed changes to various parameters, including interleaver parameter values such as interleaver depth (*see* G.992.1, § II.1.1), and that the "potential reconfiguration procedures" in discussion in the ITU-T standards group included at least bit swapping, dynamic rate repartitioning, and seamless rate adaptation. *See, e.g.*, SC-060, § 2.1. It was known in the art that both dynamic rate repartitioning and seamless rate adaptation could change PMS-TC parameters, which include interleaver parameters, such as the interleaver depth. *See, e.g.*, G.992.3, § 7.2 ("In addition to transport functionality, the ATU transmit PMS-TC function also provides procedures for: • scrambler; •

insertion of redundancy for Reed-Solomon-based forward error correction; • insertion of checksums for block based error detection; and • interleaving of data frames to spread the effect of impulsive impairments on the U interface."); *id.* at § 7.5 (listing PMS-TC control parameters, which include "D_p," defined as "[t]he interleaving depth in the latency path function #p."); *see also, e.g.*, SC-060, § 2.1 (indicating that dynamic rate repartitioning and seamless rate adaptation typically both require changes to PMS-TC parameters).

606. Accordingly, a person having ordinary skill in the art would have understood that the protocol described by Wunsch could be used by the transceivers to make "other modifications" to transceive parameters, including the interleaver depth so that a first interleaver parameter value of the first FIP setting is different than a second interleaver parameter value of the second FIP setting.

607. Accordingly, claim 10 is rendered obvious by G.992.1 in combination with Wunsch.

j. Motivation to combine G.992.1 and Wunsch

608. In my opinion, a person having ordinary skill in the art as of the priority date of the '835 patent would have been motivated to combine the disclosures of G.992.1 with the disclosures of Wunsch and would have had a reasonable expectation of success in making the combination.

609. First, Wunsch and G.992.1 are in the same field of ADSL communications, and Wunsch incorporates G.992.1 by reference. Wunsch, ¶ [0003]. In particular, Wunsch, like the '835 patent, is directed to the goal of improving on-line reconfiguration of transceiver settings in the presence of impulse noise. *See*, *e.g.*, *id.* at ¶ [0018] ("The preferred embodiment of the invention improves the ability for ADSL equipment to reconfigure transceive parameters during normal 'showtime' operation."); *id.* at ¶ [0071] ("The preferred embodiment, by employing the

the ATU-C actually receives the DRA_Swap_Reply message in which the ATU-R indicated it accepted the proposed timing of the reconfiguration. In other words, when the ATU-R accepts the proposed timing of the reconfiguration, it sends the DRA_Swap_Reply message indicating its acceptance five times, assumes that the ATU-C will receive its reply and decode it properly, and then switches to the new FIP setting(s) at the superframe corresponding to the SFR value.

613. On the other hand, the ATU-C *only* switches to the new FIP setting(s) if it receives at least three of the five DRA_Swap_Reply messages sent by the ATU-R, and it determines from those three DRA_Swap_Reply messages that the ATU-R accepted the proposed timing of the reconfiguration. Accordingly, if the ATU-C does not receive and properly decode at least three identical DRA_Swap_Reply messages indicating the ATU-R's acceptance of the proposed timing of the reconfiguration, it will continue to operate using the original FIP setting(s), whereas the ATU-R will switch to the new FIP setting(s).

614. As would have been appreciated by a person having ordinary skill in the art, and as Wunsch describes as a drawback of "current generation ADSL equipment," (Wunsch, ¶ [0013]), the reliance of the G.992.1 DRA protocol on the successful receipt by the ATU-C of three instances of the ATU-R's acknowledgment of the timing of the switch to new FIP setting(s) could be problematic in noisy channel conditions—the very conditions that would likely prompt an on-line reconfiguration in the first place. Furthermore, and as I explained above, G.992.1 introduced a reduced overhead mode in which it is possible (but not required) that only the fast buffer is used. In this case, the aoc messages used to coordinate reconfigurations during Showtime, including DRA, are transmitted in the fast path and are not protected by interleaving. As a result, when only the fast buffer is used, the reconfiguration protocol of G.992.1 is less robust (more prone to errors) than it would otherwise be.

615. Thus, a person having ordinary skill in the art would have been motivated to improve the robustness of the DRA protocol, and specifically the acknowledgment of the timing of the new configuration, to provide improved "protection against transmission errors." *See, e.g.*, BI-063, § 1 (describing need for "features that help the robustness and performance of the modems" and stating that "there is a need for a sufficiently fast, reliable and robust signaling method to allow dynamic rate repartitioning," which "can be used for bit swapping and DRA parameter changes as well and would increase the utility of these methods"). Specifically, a person having ordinary skill in the art would have sought ways to improve the robustness of the DRA protocol regardless of whether transceivers operate with dual latency or in the reduced overhead mode.

616. Wunsch recognizes the potential in "current generation ADSL equipment" for only one of the transceivers to switch to new transceiver parameters, which "causes errors in the link that degrade performance, often requiring the equipment to leave ShowTime and retrain." Wunsch, ¶ [0013]. Wunsch proposes to address this problem by improving the robustness of on-line reconfigurations during Showtime using a protocol in which a receiving transceiver sends reconfiguration transceiver parameters, implementation timing information, and acknowledgment timing information to a transmitting transceiver, and the transmitting transceiver acknowledges the proposed reconfiguration in accordance with the acknowledgment timing information. Wunsch, Abstract, Fig. 4. The acknowledgment ("ack/comply") sent by the transmitting transceiver is a Synch Flag to acknowledge and agree to implementation of the reconfiguration proposed by the receiving transceiver. *Id.* at Abstract, ¶¶ [0017], [0019]. The acknowledgment timing information instructs the transmitting transceiver when to send the Synch Flag. *Id.* at ¶ [0016].